

SANDIA REPORT

SAND2009-6008

Unlimited Release

Printed September 2009

Risk-based decision making for staggered bioterrorist attacks: Resource allocation and risk reduction in “reload” scenarios

J. Ray, P. T. Boggs, D. M. Gay, M. N. Lemaster and M. E. Ehlen

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



Risk-based decision making for staggered bioterrorist attacks: Resource allocation and risk reduction in “reload” scenarios

J. Ray, P. T. Boggs, M. N. Lemaster
Sandia National Laboratories, P.O. Box 969, Livermore, CA 94551
{jairay, ptboggs, mnlemas}@sandia.gov

D. M. Gay, M. A. Ehlen,
Sandia National Laboratories, P. O. Box 5800, Albuquerque, NM 87185
{dmgay, maehlen}@sandia.gov

Abstract

Staggered bioterrorist attacks with aerosolized pathogens on population centers present a formidable challenge to resource allocation and response planning. The response and planning will commence immediately after the detection of the first attack and with no or little information of the second attack. In this report, we outline a method by which resource allocation may be performed. It involves probabilistic reconstruction of the bioterrorist attack from partial observations of the outbreak, followed by an optimization-under-uncertainty approach to perform resource allocations. We consider both single-site and time-staggered multi-site attacks (i.e., a reload scenario) under conditions when resources (personnel and equipment which are difficult to gather and transport) are insufficient. Both communicable (plague) and non-communicable diseases (anthrax) are addressed, and we also consider cases when the data, the time-series of people reporting with symptoms, are confounded with a reporting delay. We demonstrate how our approach develops allocations profiles that have the potential to reduce the probability of an extremely adverse outcome in exchange for a more certain, but less adverse outcome. We explore the effect of placing limits on daily allocations. Further, since our method is data-driven, the resource allocation progressively improves as more data becomes available.

Acknowledgment

This work was supported Sandia National Laboratories' LDRD (Laboratory Directed Research and Development) funds, sponsored by the Homeland Security and Defense Investment Area.

Contents

1	Introduction	11
2	Literature review	15
2.1	Inference of attack parameters	15
2.2	Least-regret and multi-stage stochastic optimization techniques	16
2.3	Agent-based market simulations	18
3	Modeling reporting delay	21
3.1	Formulation for correcting the reporting delay	23
3.2	Test cases	23
4	Resource allocation in attacks with non-communicable agents	31
4.1	Formulation	31
4.1.1	Optimal allocation of resources	32
4.2	An attack on one city	37
4.3	An attack on two cities: The “reload” case and the equilibration of pain	38
4.4	An attack on one city with corrected data	40
5	Resource allocation for outbreaks caused by communicable diseases	47
5.1	Disease dynamics	48
5.2	Formulation of the inverse problem	48
5.3	Test case 1: A plague outbreak	50
5.3.1	The outbreak simulation	50
5.3.2	Inference of outbreak parameters	51

5.4	Test case 2: An influenza outbreak	53
5.4.1	The outbreak simulation	56
5.4.2	Inference of outbreak parameters	57
6	Conclusions	61
	References	63

List of Figures

3.1	Number of individuals who have reported by Day $R = 6$ having developed symptoms on days $i = [0, 5]$. The long reporting delay results in a large discrepancy between the number of individuals who have sought treatment by Day 6 (blue) and the number who actually developed symptoms on a given day (red). Also shown is our estimation of the total number to develop symptoms (black), which closely tracks the red curve except for the final point.	26
3.2	Comparison of the 1D PDFs of the number of individuals to be infected (top), the time of infection (middle), and the logarithm of dosage (bottom), based on the pristine data and the corrected data on Day $R = 6$, as presented in Figure 3.1.	27
3.3	Evolution of the epidemic using 4000 samples taken from the joint PDF of the attack parameters based on the corrected data on Day $R = 6$ as presented in Figure 3.1. The actual number of people showing symptoms by Day 6 is plotted using symbols.	28
3.4	Evolution of the epidemic using 4000 samples taken from the joint PDF of the attack parameters based on the pristine data, for comparison to Figure 3.3. The actual number of people showing symptoms by Day 6 is plotted using symbols.	29
4.1	Multi-scenario resource allocation problem.	35
4.2	Allocations for the attacked city, obtained from data collected over the first 5 days. The gray region denotes the evolution of all the scenarios considered. The net effect of σ is to reduce the allocation during the early days of the epidemic. $R = 10,000$. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds.	41
4.3	PDFs of excess casualties for $\sigma = 0.04$ and 0.1 . $R = 10,000$. A tighter daily constraint on allocations ($\sigma = 0.04$) increases the probability of excess casualties. However, note that the PDFs have rather compact support.	42
4.4	Allocation under various values of σ for City A. The gray region denotes the evolution of all the scenarios considered. $R = 50,000$ (total for both cities). These allocations were drawn from data collected over 5 days; allocations are for Day 6. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds.	43

4.5	Allocation under various values of σ for City B. $R = 50,000$ (total for both cities). These allocations were drawn from data collected over 3 days (Days 3, 4 and 5, the attack stagger being 2 days); allocations are for Day 6. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds. The gray region denotes the evolution of all the scenarios considered. . .	44
4.6	PDF of excess casualties for $\sigma = 0.02, 0.04$ and 0.1 . $R = 50,000$. Note how $\sigma = 0.1$ results in certain “negative” casualties, i.e., excess resources in certain low casualty scenarios. Most of the excesses occur in the smaller attack (City A). The PDF corresponding to $\sigma = 0.02$ and 0.04 have a similar support widths but the higher value of σ reduced the expected value of casualties.	45
4.7	Resource allocation profiles over the durations Days 4-9 calculated from inferences drawn from increasing amounts of data. We see the effect of the 3-day ramp up. Further, most of the allocation is done early with a quick curtailing of allocation later in the outbreak.	46
5.1	Evolution of the plague epidemic. New cases of symptomatic patients are plotted in red while the latent infection rate (daily infections) are in blue. The simulation was started with 100 index cases distributed in the social network. We see, from the symptomatic cases, that the epidemic peaks around Day 15. The infection rate peaks a few days before.	52
5.2	The MCMC chains for the outbreak parameters $\{N_{tot}, \alpha, \tau, k\}$, plotted in sequence, from top to bottom, on the left. On the right are the histograms of the samples of the outbreak parameters. These were developed from a time-series 20 days long. The figures for θ were left out for lack of space.	54
5.3	The latent infection intensity curve, developed from the maximum likelihood estimates of the outbreak parameters, using 10 and 25 days of data. The true infection intensity is also plotted. The blue dot at the left extreme indicates the index cases; we have not plotted the numerical estimate derived from the time-series.	55
5.4	Evolution of the influenza epidemic. New cases of symptomatic patients are plotted in red while the latent infection rate (daily infections) are in blue. The simulation was started with 100 index cases distributed in the social network. We see that both the time-series exhibit an upward trend.	58

List of Tables

3.1	Number of individuals, $N_{i,r}$, turning symptomatic on day i as known by day r	22
3.2	Ratios $N_{i,r}/N_{i,r+1}$ computed for fitting the fraction of individuals that turned symptomatic on day i that have sought treatment within $\Delta = r - i$ days.	24
3.3	Each column provides the estimate M_i of the number of individuals turning symptomatic on day i , using only data collected on days $r < R$. The final column shows the actual number of symptomatics, M_i^* , for comparison.	25
4.1	Effect of ρ on a and expected deaths (4.7).	37
5.1	Estimates (medians) of the outbreak parameters developed from time-series of different lengths. The 95% confidence intervals are mentioned in parenthesis and the true values are in square brackets.	53
5.2	Summary of influenza characterization from a variety of sources. All estimates of duration are in days.	57
5.3	Summary of parameters for our influenza model.	58
5.4	Estimates (medians) of the outbreak parameters developed from time-series of different lengths. The 95% confidence intervals are mentioned in parenthesis and the true values are in square brackets.	59

This page intentionally left blank.

Chapter 1

Introduction

This report describes a set of resource allocation techniques developed to address the problem of a “reload” scenario, i.e., a series of time-staggered bioterrorist attacks conducted over population centers with an aerosolized pathogen. Such a problem is extremely challenging since it requires one to allocate resources early, after the detection of the first attack and in ignorance (or with little knowledge) of the subsequent attacks. Resources here refer to those which are difficult to gather and transport and may thus be considered scarce, for example, medical equipment and personnel. Current approaches leave much to be desired, since they make scant use of the only source of data in the aftermath of a successful attack, i.e., the morbidity stream, which typically consists of the number of people showing symptoms. Instead, they rely heavily on detection via aerosol sensors. Such an approach has its merits; early detection, followed by heavy prophylaxis, has the potential to prevent an outbreak [1]. However, if the aerosolized pathogen is not detected (for example, if the attacked site is not instrumented with sensors), an outbreak may be expected, leading to a rapid and large increase in demand for medical resources. A proper resource allocation technique would consist of drawing estimates of the resource demand from the available data, and performing the resource allocation accordingly. Since the data stream may reasonably be expected to become more informative as one progresses into the outbreak, demand estimates and resource allocations should be dynamically updated for accuracy. In this report, we demonstrate how the estimation may be performed probabilistically, followed by a resource allocation using the uncertain resource demand estimates. Our formulation allows the calculation of the risk associated with the allocation, and is a function of a free parameter, representing the risk appetite of the resource allocator or alternatively, the risk associated due to the constraints of the transportation infrastructure. The results here are preliminary, but show promise for both non-communicable and communicable diseases. Further, the techniques are not specific to bioattacks; they can be used in other scenarios where the effects of an attack are delayed (i.e., there is an “incubation” time) and an estimation has to be performed to determine the cause. Thus the hazards posed by the percutaneous absorption of chemical agents (or toxic industrial chemicals) which have an “incubation” period of about a day, as well as radiation exposure (somewhat larger delays) are also problems where our resource allocation technique may find use.

The problem of the “reload” scenario was first described by Danzig [2]. The report had a strong policy focus, and did not delve into technical means of addressing the problems that could be expected in a “reload” scenario. A detailed system-dynamical study was performed by Edwards *et al.* [3] and certain technical and procedural changes/targets were suggested. However, the emphasis was on detection rather than response. The problem of an undetected attack (or multiple small

undetected attacks) was not considered.

The problem of resource allocation in case of an undetected bioattack presents some interesting mathematical challenges. In such a context, the only source of data is the time-series of patients infected with the aerosolized pathogen in question. This information stream will be embedded in the normal background morbidity stream that might be expected in any population center. During the aftermath of the bioattack, the background morbidity will be augmented by hypochondriacs/“worried-well” individuals. Further, time-series of infected people showing symptoms will be confounded with a random reporting delay – infected, symptomatic people cannot be expected to report to medical institutions immediately after the exhibition of symptoms. Thus the “signal” of the attack in the morbidity time-series may suffer from a low signal-to-noise ratio for a significant duration, before the anomalously large morbidity level due to the attack triggers an alarm.

In the absence of a background morbidity, the “signal” from a bioattack will consist of the time-series of infected patients turning symptomatic and reporting to medical institutions for care. The delay between infection and exhibition of symptoms is the incubation period, and is characteristic of a disease, i.e., if the etiological agent is known, a model for the incubation period is generally available. The delay between the exhibition of symptoms and reporting for care is called the “reporting delay” in this report. It has been modeled for particular populations (usually as a log-normal distribution), but will probably differ from site to site (i.e., it is best extracted from the data at hand, with perhaps existing reporting delay models to guide us). The signal will also depend on the number of people infected, N (and in case of communicable diseases, this number can be further divided into the index cases N_{ind} , a stationary number, and secondary infections N_{sec} , a time-variable quantity), the time of infection τ and in case of dose-dependent incubation periods (e.g., in anthrax), the dose D . Thus the time-series could be used to back-calculate $\{N, \tau, \log_{10}(D)\}$ which we will henceforth refer to as the *attack parameters*. Since the inference will need to be performed early after the detection, the time-series will be short (and noisy), rendering any estimate of the attack parameters uncertain; thus they are best inferred as probability density functions (PDFs). Samples of attack parameters, drawn from the PDFs conditioned on time-series data could be used to perform posterior predictive runs (with a conventional epidemic model) to bound the possible evolution of the outbreak. The ensemble of runs embodies a large fraction (but not all) of the uncertainty regarding future demands on resources (e.g., it would not capture the uncertainty / inaccuracies in the epidemic model), and allows one to consider a resource allocation procedure that could reduce the chances of an extremely bad outcome. Defining what a desirable outcome might be is case dependent, but will extend beyond minimizing the expected number of casualties and will be addressed in this report.

The question of an “optimal” allocation of resources, given an uncertain demand for them, is addressed using a multistage optimization technique. In case of a “reload”, one will have uncertain demands at multiple sites, with differing levels of uncertainty. However, both the epidemic and the demand for resources will evolve over time, allowing one to consider a *time profile* of resource allocation, which can be updated as the morbidity time-series lengthens. The updating will be constrained by the capacity of the transportation infrastructure as well as the risk-appetite of the emergency manager, if subsequent attacks are expected and resources have to be husbanded. The

resource allocation profile can be used to dispatch resources in the short term (timescale of a day) and plan for transportation needs in the long term ($O(\tau)$ days). The constraints placed on the resource allocation profile can have a significant impact on the risk (henceforth defined as the PDF of casualties, given an ensemble of outbreak realizations and a resource allocation profile). These constraints can be parameterized and we investigate the sensitivity of the risk to these free parameters. The parameterized constraints are termed “free” since they are governed by transportation and/or risk appetite, which are exogenous to the information content of the morbidity time-series and thus cannot be informed by better data analysis.

Apart from a multistage optimization technique, resource allocation may also be performed using a “market-based” technique, i.e., a system where individual entities (“agents”) negotiate/converge to a resource allocation profile depending upon their need and supplies of resources. Such a system requires the relevant information (on the need for resources at various sites) to permeate throughout the system (the “market”) via interactions between “agents” (one of whom may be a “market-maker”). This is a novel technique, but has not been explored in this study. This has the potential to arrive at better resource allocation in real life, but faces many modeling and computational challenges. However, we include a review of literature on this area in this report.

The report is structured as follows. In Chapter 2, we review recent work on the inference of attack/outbreak parameters from time-series of morbidity data. We also include a short discussion of the reporting delay, as well as a review of robust techniques in multi-stage stochastic optimization. The chapter also contains a description of “market-based” and “agent-based” modeling approaches to resource allocation. In Chp. 3 we address the question of how the reporting delay may be estimated and used to correct the morbidity time-series (which consist of symptomatic patients who have reported to medical institutions) to obtain an approximation the actual number of infected, symptomatic people (some of whom may not have begun to seek medical attention). The impact of the “corrected” time-series on the inference of attack parameters will be studied. In Chp. 4 we formulate the resource allocation problem (including the means to accommodate multi-site demands with disparate levels of uncertainties in the demand estimates) and present examples using a non-communicable disease, specifically anthrax. We study the effect of the “free” constraint parameter, and also show how a “naive” allocation (an allocation based on the mean of the ensemble of outbreak realizations) leads to a very risky allocation. In Chp. 5 we show how the inference procedure can be extended to address a communicable disease, specifically, H1N1 flu modeled on the 1918 pandemic; the process of conducting posterior predictive runs and performing the resource allocation is the same as in anthrax and is omitted. In Chp. 6, we summarize our findings, and identify topics which need further investigation. These may be needed to accurately gauge the potential of our technique for responding to both bioattacks and epidemics/pandemics engendered the increased contact between humans and wildlife, spurred by both economic, climate change and spread by globalization.

This page intentionally left blank.

Chapter 2

Literature review

In this chapter, we review some existing literature on the inference of outbreak/attack parameters, stochastic, multi-stage optimization and agent-based market simulations. The latter are two different radically different approaches to computing resource allocations,

2.1 Inference of attack parameters

The approaches used to characterize (i.e., estimate attack parameters) partially observed outbreaks are very different for non-communicable and communicable diseases. For non-communicable diseases, current literature consists of a few studies involving releases of aerosolized anthrax. Walden & Kaplan [4] introduced a Bayesian formulation for estimating the size and time of a bioterror (BT) attack and tested it on a low-dose (less than ID_{25} , the dose at which a person has a 25% probability of incurring the disease) anthrax release corresponding, approximately, to the Sverdlovsk outbreak [5] of 1979. Their formulation incorporated an incubation period model developed by Brookmeyer *et al.* [6] and demonstrated the use of prior distributions on N to reduce uncertainty in the inferred characteristics. Brookmeyer & Blades [7] used a maximum likelihood approach, along with the anthrax incubation model in [6], to infer the size of the 2001 anthrax attacks [8] before estimating the reduction in casualties due to the timely administration of antibiotics. Both [4] and [7] developed similar expressions for the likelihood function, i.e., the probability of observing a patient time series given an attack at time τ with N infected people. The incubation period model in [6] was not dose-dependent, and hence no doses were inferred in these two studies.

The BARD [9] effort also seeks to characterize a BT attack from the presentation of symptoms. It attempts to estimate the location, height, and time of an airborne anthrax release, as well as the number of spores. The observables consist of respiratory visits to emergency departments, as might be obtainable from syndromic surveillance systems such as RODS [10]. The model that relates these observables to the characteristics of the outbreak includes a Gaussian dispersion plume [11], Glassman’s infection relation [12], and a log-normal distribution of incubation periods, with dose-dependent mean and standard deviation. However, BARD’s use in an urban context is only approximate since Gaussian plumes are suited mainly for open spaces [11]. Further, estimation of the release parameters was an intermediate aim in [9]; its thrust is to detect anomalous morbidity patterns using a spatiotemporal approach, which is considerably assisted if the spatial distribution of infected people (i.e., the “footprint” of the plume) can be estimated. A similar,

spatiotemporal approach to attack parameter estimation can be found in Legrand *et al.* [13]. The formulation is Bayesian, and the paper contains a thorough testing of their estimation technique against competing ones. The study also dealt with how medical resources/care could be prioritized spatially, based on the severity of infection in different locations in the attacked site.

In this report, the inference of attack parameter is performed using the Bayesian technique described in [14]. The method is solely temporal and thus has simple data needs (it does not need any spatial information, unlike [9] and [13]). It yields estimates of the attack parameters $\{N, \tau, \log_{10}(D)\}$. The data consists of a time-series of the number of new patients exhibiting symptoms on a daily basis. The model validation performed in [14] demonstrated that about 5 days of data are sufficient to develop informative PDFs of the attack parameters and perform posterior predictive runs. In certain cases, 5 days of data could lead to a wrong estimation, but they were quickly corrected as more data became available. However, a shortcoming of the technique is that it requires a time-series of the *actual* number of symptomatic individuals, including those who had not reported for medical care. In this study we will investigate how such a time-series may be approximated from a time-series of symptomatic individuals who seeking medical care. The difference between the two arises from a reporting delay which has been modeled [15], and which can be explicitly estimated from data. This is described in Chp. 3.

The spread of communicable diseases shows many dynamical features and gives rise to a different parameter estimation problem. Traditionally, this has meant estimating the rate of spread of a disease from data. There is a vast literature on fitting conventional SEIR models to data [16, 17, 18, 19, 20, 21], and of late, this has been extended to network-based epidemic models [22] as well as inferring chains of transmission [23]. A simple approach to estimating outbreak parameters is described in [24]. The authors assume that there exists a time-dependent infection intensity and an unknown number of infected and infectious individuals. The (unknown) infection intensity, convolved with the incubation period of the disease results in the time-series of people exhibiting symptoms. This technique was used to back-calculate the number of individual infected with HIV using the data collected in 1980–1988. In this work, we will use this simpler model of a communicable disease to infer the number of secondary infection as well as estimate the shape of the infection intensity curve.

2.2 Least-regret and multi-stage stochastic optimization techniques

Given that we can obtain a probability density function that captures the uncertainty in the extent of an anthrax attack, the problem becomes one of optimally allocating resources to minimize the expected number of deaths. The idea is to sample the PDF to obtain scenarios of the number of people who will arrive at the hospital each day over the extent of the attack. As we expected, these scenarios can, and do, vary widely, especially early in the attack when little is known from which to calculate the PDFs. The challenge is to create an optimization model, based on these scenarios, that takes into account other relevant constraints, including logistics limitations, including the risk appetite of the emergency manager, and social, or fairness, constraints that allocate resources proportionally over several attack sites. We give two approaches, a so-called “least-regret” model and

a more classic stochastic optimization with recourse model. We combine these to get an optimal allocation for today from the stochastic optimization model along with a proposed allocation over the entire period from the least-regret model.

The idea behind the least-regret approach is first to solve a resource allocation problem for each of the scenarios. Since for each scenario we know exactly how many people arrive and we know the expected effectiveness of our resources, the optimization problem is a classical resource allocation problem that is easily solved. In fact, such problems are straightforward linear programming problems; many implementations of fast, reliable algorithms are available to solve them. After solving these problems, we have a resource allocation schedule and the minimal number of deaths for each scenario. At this point we could make a naive choice of simply using the average of all of the allocations for today as our choice, but this has certain drawbacks, as we report in [25]. A better strategy is to solve a final optimization problem that picks an allocation schedule such that we deviate as little as possible from the minimal number of deaths in each of the scenarios. Specifically, let r be an allocation schedule, i.e., r_j is the allocation made on day j . Let D_i^* be the minimum number of deaths obtained in scenario i and let $D_i(r)$ be the number of deaths that would occur if allocation r is used in scenario i . Then the least-regret problem is to minimize over all allocation schedules r the quantity

$$\sum_{i=1}^K (D_i(r) - D_i^*)^2,$$

where K is the number of scenarios. This allocation has the advantage of not allowing any scenario to dominate the calculation. Of course, other criteria could be postulated, but this formulation is appealing on both computational and practical grounds. From a computational point of view, this will result in a quadratic programming problem for which there are many good algorithms.

In our experimentation and testing we included a number of constraints that are meant to show the ability of the model to handle a variety of situations that may reasonably arise. For example, we considered bounds on the number of units of resource that could be shipped on each day, we included a “ramp-up” at the beginning to allow for a logistics resources to be put in place, we included an assumption that allowed for the effectiveness of treatment to vary over the course of the attack, and we allowed for the reuse of resources if patients receiving them died.

We also included in our model the possibility of an attack on another target (the “re-load” case) within a few days of the first attack. In this case we had to explicitly constrain the solution to make allocations to both cities. Without this, the optimization model could achieve the same minimal number of deaths by allocating all (or most) resources to only one city. Our assumption is that such an allocation would be socially unacceptable, so we added a constraint that ensured that all targets received proportional allocations.

We now describe the stochastic optimization model, which includes the above constraints.

Stochastic optimization with recourse provides another, more conventional way to choose resource allocations under uncertainty. If we have some idea about the probabilities of events in subsequent

days, perhaps (as in the present application) based on historical data and incoming observations, we can formulate models that account for corrective actions — recourse — we can take once more is known, and we can make today’s decisions in light of expected recourse costs.

With both least-regret and stochastic programming formulations, the goal is to make a reasonable decision about what resources to allocate today. Tomorrow we will have gathered more data and can rerun the calculations with updated data to help make tomorrow’s allocation decision. In our recourse formulation, we used today’s allocation as the first stage and allocations for the subsequent days under each of the scenarios generated from today’s data as the second stage, and with expected deaths as the cost of recourse. That is, our objective was to minimize expected deaths, subject to the previously described constraints. This results in a single linear programming problem (a “deterministic equivalent”) that is larger than least-regret problem, but faster to solve. Both formulations give similar decisions about today’s allocation. A detailed description of our recourse formulation appears in [25] and in Chapter 4 below.

2.3 Agent-based market simulations

The previous sections have dealt with centrally planned allocation of resources (goods) in response to a crisis. Allocation of goods can also be performed using market mechanisms, based on a price of the good, attributes of the buyers and sellers or government policies. Over the last 15 years, agent-based simulations have been used to model and understand how markets function, where a market in the formal sense is “any context in which sale and purchase of goods and services takes place” [26]. Each market tends to be for a single good (e.g., vaccines), sellers are those who are willing and able to sell the particular good and buyers are those willing and able to purchase it. As compared with the centrally planned allocations of these goods and services, markets provide mechanisms for allocating based on the price of the particular good or service. Given that in a real market individual buyers and sellers have private encapsulated knowledge about themselves (as well as about other buyers and sellers) and have private encapsulated procedures that they follow and that the resulting market behavior based on the public actions of these buyers and sellers is dynamic, nonlinear and complex, agent-based models are an ideal approach.

Agent-based market simulations have been used extensively to model a wide array of markets, starting largely from the seminal work of Palmer et al [27], Axelrod [28], Arthur et al [29] and Epstein and Axtell [30], the later of whom formalized the notions of agents as “people” in “environments” that follow “rules”, all of which results in emergent “social structures”. Their work indicated that the traditional theories of market equilibrium were very sensitive to the set of market conditions (buyer behavior, seller behavior, numbers of buyers and sellers, information and market-clearing mechanisms).

Since then, there have been a vast expansion of research in many market domains that focus on analyzing market design (auctioning and other market mechanisms for rationalizing resources from sellers to buyers) and learning and adaptive seller/buyer behaviors. A prominent example is agent-based simulation of spot markets for wholesale electric power, where wholesale electric power that

has not already been sold through long-term contracts is sold in short-term markets. As described in [31], many electric power market models have been developed primarily to analyze power buyer and seller learning and resulting market dynamics, complexity and ultimately the reliability of these spot markets to deliver power and efficient prices. A significant motivation for this body of research was the failure of electric power spot markets in California, where poor spot market designs resulted in market-induced supranormal prices and profits and ultimately rolling “brown-outs”.

Another prominent example is agent-based simulations of financial markets, where changes in market design and buyer/seller behaviors can have complex, unforeseen outcomes in market prices, i.e., the value of assets traded and stability i.e, the liquidity and fluctuations in valuations that can destabilize and even stop a market from functioning. As described in [32], NASDAQ has used agent-based simulations to analyze impact of regulatory changes on their market under various changes in buyer/seller strategies, price increments and so on; furthermore, eBay uses intelligent agents to help its customers with their market bidding.

Ehlen et al [33] used Sandia N-ABLE model to analyze how real-time pricing of consumer power would affect wholesale and transmission market pricing and stability; Sprigg and Ehlen [34] use ASPEN [35] to simulate how agents can find their Nash equilibrium prices with little information and simplistic decision rules. Agent-based simulations have also been used to investigate how spot, future and option markets could be destabilized by terrorist events [36].

This page intentionally left blank.

Chapter 3

Modeling reporting delay

As discussed in Chp. 2, Sec. 2.1, the estimation of the demand for resources requires one to characterize the number of index cases, and if the disease is communicable, the secondary cases too. This estimation is performed using the time-series of new people exhibiting symptoms. However, since symptomatic people do not seek care immediately on exhibiting symptoms, this time-series is frequently not available. What is available is a time-series of symptomatic patients who have sought care. Given a daily time-series of length R , the reporting delay (which, in case of the Sverdlovsk anthrax outbreak was characterized as a lognormal distribution with a median of 2.7 days [15]) leads to a severe under-reporting of the number of symptomatics over the time interval $[R - 5, R)$. Any inferences drawn with such an erroneous time-series will be misleading; ignoring the last 5 days of data is not an option since that would increase the length of the observation period. In a situation where a response has to be formulated quickly, such an approach would be useless. Thus one has to consider either “correcting” the data for the reporting delay, or formulating an estimation problem for the index cases which incorporates the delay.

The number of people seeking care on a given day can be always be asked about their time of exhibition of symptoms. Thus for a given reporting day r , one can construct a time-series $N_{i,r}$ consisting of the number of people who exhibited symptoms on day $i, i \leq r$. One can curate such time-series over R days to obtain a table. An example of such a table is Table 3.1, which was generated from a simulated anthrax attack (described in Sec. 3.2), with a reporting delay model obtained from [15]. Here, each column r contains information on the number of people showing symptoms on each day $i, i < r$. The rows of the table denote the day i .

Compare the column $r = 6$ with the last column, $r = 9$. $N_{0,6}$, the number of people turning symptomatic on day 0, as known on reporting day 6 (alternatively, an approximation of M_0 , the number of people who turned symptomatic on day 0) is not very different from the estimate $N_{0,7}$. This is because most of M_0 symptomatic patients have sought medical care (i.e., have been reported) by Day 6. However, if one considers the case of M_6 , $N_{6,6}$ and $N_{6,9}$ differ by an order of magnitude. In fact, each row of the table traces out the cumulative distribution function of the reporting delay, and asymptotes to the true number of M_i of the number of index cases who turn symptomatic on day i . If one assumes a model for the reporting delay distribution (e.g., a log-normal or Γ -distribution), one can obtain the model parameters by regressing to the data in the table; obtaining the asymptotic value M_i , given the first few days of observations $N_{i,r}, r < R$ is then a trivial exercise.

An alternative approach to estimating the attack parameters is to incorporate the reporting delay in the model used for estimation. The data consists of the new number of symptomatic cases seeking

Table 3.1. Number of individuals, $N_{i,r}$, turning symptomatic on day i as known by day r .

	$r=0$	1	2	3	4	5	6	7	8	9
$i=0$	0	16	45	70	89	100	108	115	116	117
1	-	14	161	328	457	548	605	649	678	700
2	-	-	31	338	710	1006	1207	1340	1415	1481
3	-	-	-	44	509	1077	1484	1759	1932	2057
4	-	-	-	-	65	570	1211	1669	2008	2219
5	-	-	-	-	-	48	453	1042	1540	1835
6	-	-	-	-	-	-	62	497	1053	1501
7	-	-	-	-	-	-	-	51	428	880
8	-	-	-	-	-	-	-	-	39	379

care, collated on a daily basis. This data can be obtained by performing column-wise sums in Table 3.1 (which would provide a running sum of the total number of people who have sought care by day r) and subtracting the column-wise sums from their predecessor (which would provide the new daily cases). Consider that there exists an unknown number N of index cases, who were infected τ days before the reporting of first symptoms. Assume that an average dose D was the infecting dose. Let $f_I(x; D)$ represent the dose-dependent incubation period of anthrax. Then the number of people n_i seeking care on day i , i.e., in the time interval $[t_{i-1}, t_i]$ is given by

$$n_i = N \int_{t_{i-1}}^{t_i} f_I(s; D) (C(t_i - s; \mathbf{p}) - C(t_{i-1} - s; \mathbf{p})) ds, \quad (3.1)$$

where $C(t; \mathbf{p})$ is the cumulative distribution function (CDF) of the reporting delay. This formulation was adapted from [24]. The parameters of the CDF, \mathbf{p} , along with N, τ and D , can be estimated from the time-series n_i .

In this work, we will investigate the first approach rather than the one based on Eq. 3.1. The efficacy of Eq. 3.1 has been investigated in [24]; furthermore, it estimates more parameters (N, τ, D, \mathbf{p}) from less data (one time-series) than the tabulation approach that regresses a distribution model to a table of data. Furthermore, it allows one to compare the efficacy and applicability of validated models that exist in literature with a more realistic data-stream.

Below, we present a formulation that allows us to model the reporting delay as Γ -distribution. The parameters of the distribution are estimated from the data in Table 3.1 using a least-squares method, and thereafter used to correct the observation to obtain estimates of M_i . We then use M_i to infer the attack parameters, and compare them with inferences drawn from M_i^* , the true time-series of symptomatics. These tests are performed using synthetic data from a simulated anthrax attack. The reporting delay in the simulated attacks is modeled using the log-normal distribution in [15]. We also compare the posterior predictive runs obtained from M_i and M_i^* to gauge what the impact of reporting delay correction might be on the allocation mechanism.

3.1 Formulation for correcting the reporting delay

In correcting for the reporting delay, we assume that the PDF of the delay between when an individual turns symptomatic and seeks treatment will follow a Γ -distribution. The fraction $f_{i,r} = N_{i,r}/M_i^*$ of people who turned symptomatic on day i as known by day r will then follow the Γ -distribution CDF,

$$F(x; k, \theta) \equiv \frac{\gamma(k, x/\theta)}{\Gamma(k)} = \frac{\int_0^{x/\theta} q^{k-1} e^{-q} dq}{\int_0^\infty q^{k-1} e^{-q} dq}, \quad (3.2)$$

with shape parameter $k > 0$ and scale parameter $\theta > 0$ fixed over time. Here, $\Gamma(k)$ is the Gamma function and $\gamma(k, x/\theta)$ is the lower incomplete Gamma function. This implies that the CDF of the reporting delay depends only on the delay between turning symptomatic and seeking treatment, $\Delta = r - i$. Due to the coarse binning of the data, we allow a shift between Δ and x , using $x = \Delta + \delta$, where δ is our third fitting parameter.

Because we do not know M_i^* , however, we will instead fit the ratio $f_{i,r}/f_{i,r+1} = N_{i,r}/N_{i,r+1}$ (illustrated in Table 3.2 for the test case), in which the unknown term cancels, using the fitting function

$$\eta(\Delta; k, \theta, \delta) \equiv \frac{F(\Delta + \delta; k, \theta)}{F(\Delta + \delta + 1; k, \theta)}. \quad (3.3)$$

The regression was performed using ITT Visual Information Solutions IDL with the routine MPFIT by Craig Markwardt. MPFIT is based on the MINPACK-1 Fortran package for least-squares minimization. When performing the regressions described below, we omitted the day $i = 0$ as well as the delay $\Delta = 0$ data. After the CDF model parameters (k, θ, δ) are determined from the regression, estimates of M_i^* (given in Table 3.3) are computed using $M_i = N_{i,R}/f_{i,R}$, where R is the most recent reporting day.

3.2 Test cases

A simulated population of 80,000 people was exposed to Anthrax, resulting in 23,917 being infected. The average dose of those infected was 2,754 spores. The elapsed time between developing symptoms and seeking treatment for each individual was drawn from a log-normal distribution,

$$f(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma x} \exp\left[-\frac{\log(x/\mu)}{2\sigma^2}\right], \quad (3.4)$$

with $\mu = 2.73$ and $\sigma = 0.7$, consistent with the Sverdlovsk incident, where x , μ , and σ are given in days. Individuals turning symptomatic on days $i = [0, 8]$ who sought treatment by day $r = 9$ were included in the observations (see Table 3.1).

Table 3.2. Ratios $N_{i,r}/N_{i,r+1}$ computed for fitting the fraction of individuals that turned symptomatic on day i that have sought treatment within $\Delta = r - i$ days.

	$\Delta=0$	1	2	3	4	5	6	7	8
$i=0$	0.0000	0.3556	0.6429	0.7865	0.8900	0.9259	0.9391	0.9914	0.9915
1	0.0870	0.4909	0.7177	0.8339	0.9058	0.9322	0.9572	0.9686	-
2	0.0917	0.4761	0.7058	0.8335	0.9007	0.9470	0.9554	-	-
3	0.0864	0.4726	0.7257	0.8437	0.9105	0.9392	-	-	-
4	0.1140	0.4707	0.7256	0.8312	0.9049	-	-	-	-
5	0.1060	0.4347	0.6766	0.8392	-	-	-	-	-
6	0.1247	0.4720	0.7015	-	-	-	-	-	-
7	0.1192	0.4864	-	-	-	-	-	-	-
8	0.1029	-	-	-	-	-	-	-	-

Each column of Table 3.3 shows the corrected values as determined using data available only up to day $R = r$. Simply for comparison, the final column shows the actual number of individuals who developed symptoms on a given day. This pristine data would not be available in the case of a real attack but is useful in analyzing the accuracy of our correction as applied to the simulated data. For reporting day $R = 6$, we will present PDFs drawn from the corrected data as compared to those drawn from the pristine data. Figure 3.1 directly compares the raw observations, the observations after being corrected for the reporting delay, and the pristine data for the case being analyzed. The corrected curve follows the pristine data closely with the exception of on the most recent day. Our goal here is to determine how sensitive the PDFs will be to differences between the corrected and pristine data.

Figure 3.2 compares the 1D PDFs of the number infected, the time of infection, and the logarithm of dosage, as determined from the corrected observations and the pristine data. On Day $R = 6$, the pristine data shows a bimodal distribution for the number infected, shown in the top panel, with the primary peak most closely reflecting the actual number infected. The PDF developed from the corrected data is much more narrow and suggests a significantly smaller infected population. If we take an additional day of data and perform this comparison for $R = 7$, we find that the PDF developed from the pristine data is no longer bimodal and is better matched by the PDF from the corrected data. For $R = [8, 9]$, we find that the PDF from the corrected data very closely matches that from the pristine data, showing that our correction should provide very reliable results as the epidemic progresses.

The middle panel of Figure 3.2 shows the PDFs of the time of infection. We find that the peak for the corrected data is close to that for the pristine data when $R = 6$. When looking at $R = [7, 9]$, we find that the PDFs match for the corrected and pristine data, again validating our method. In the bottom panel, we see that the PDF of the logarithm of dosage is much wider for the pristine data as compared to the corrected data for $R = 6$. As was the case for the number infected, the difference between the PDFs for dosage is much smaller for $R = [8, 9]$.

Table 3.3. Each column provides the estimate M_i of the number of individuals turning symptomatic on day i , using only data collected on days $r < R$. The final column shows the actual number of symptomatics, M_i^* , for comparison.

	$R=4$	5	6	7	8	9	M_i^*
$i=0$	137	115	118	121	119	119	123
1	889	700	708	707	714	725	753
2	2051	1541	1567	1549	1545	1569	1618
3	3323	2318	2305	2249	2240	2263	2340
4	-	2557	2626	2549	2580	2598	2651
5	-	-	2058	2226	2371	2385	2477
6	-	-	-	2256	2278	2340	2350
7	-	-	-	-	1973	1926	1998
8	-	-	-	-	-	1760	1690

Figures 3.3 and 3.4 show 4000 samples from the joint PDF of the attack parameters as determined from the corrected observations and the pristine data, respectively. As was seen in the 1D PDFs, the distribution is significantly narrower for the corrected data as compared to the pristine data. This could give an inaccurate picture of the uncertainty in the attack parameters when determining the resource allocations. However, the corrected data give us a starting point for making resource allocations that was not available with only the raw data.

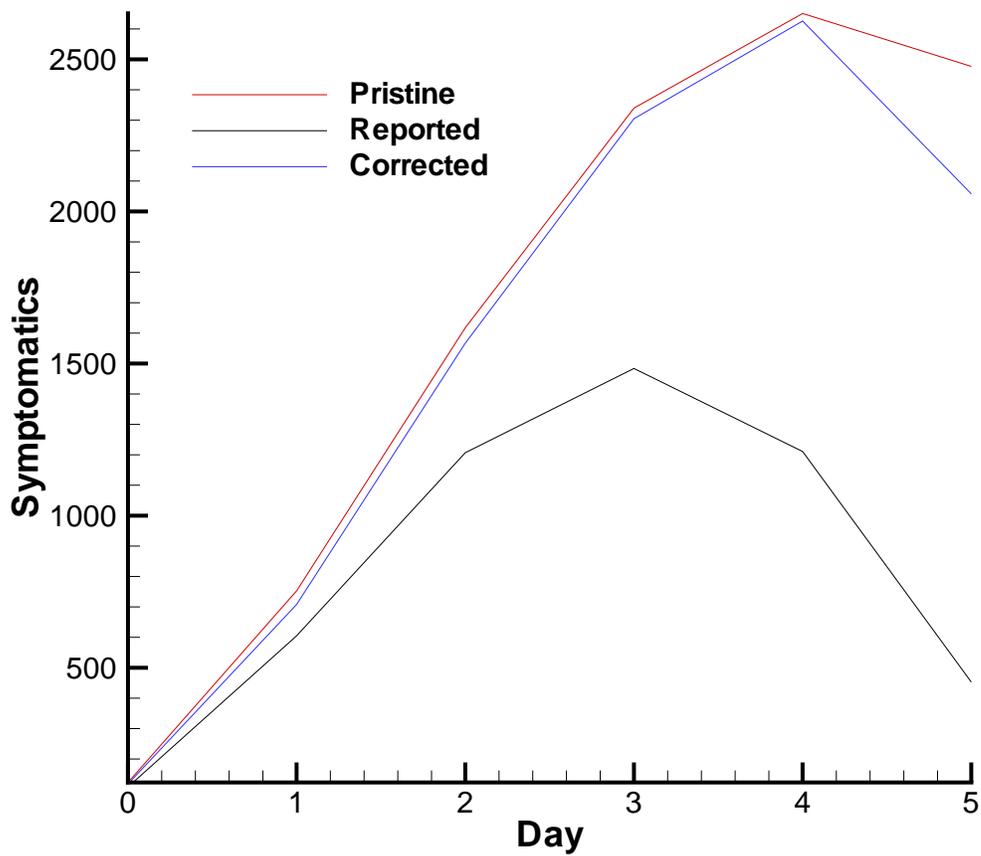


Figure 3.1. Number of individuals who have reported by Day $R = 6$ having developed symptoms on days $i = [0, 5]$. The long reporting delay results in a large discrepancy between the number of individuals who have sought treatment by Day 6 (blue) and the number who actually developed symptoms on a given day (red). Also shown is our estimation of the total number to develop symptoms (black), which closely tracks the red curve except for the final point.

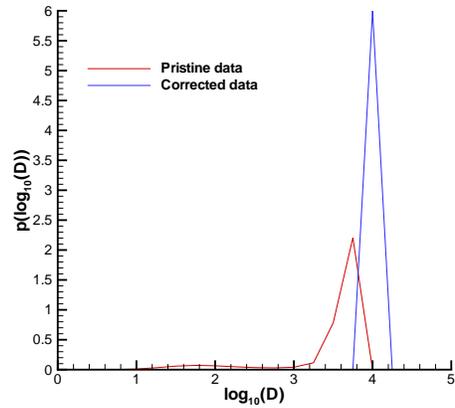
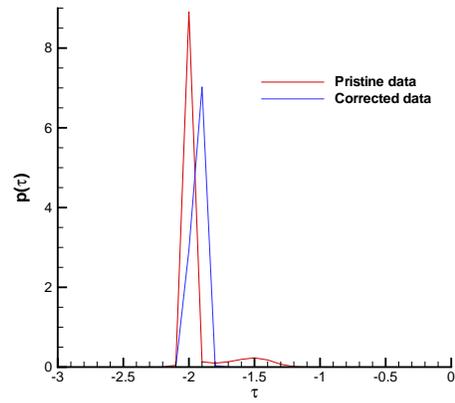
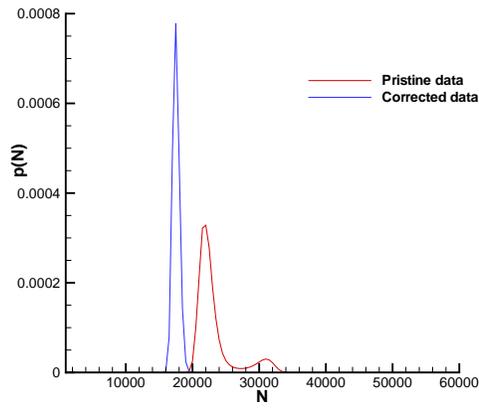


Figure 3.2. Comparison of the 1D PDFs of the number of individuals to be infected (top), the time of infection (middle), and the logarithm of dosage (bottom), based on the pristine data and the corrected data on Day $R = 6$, as presented in Figure 3.1.

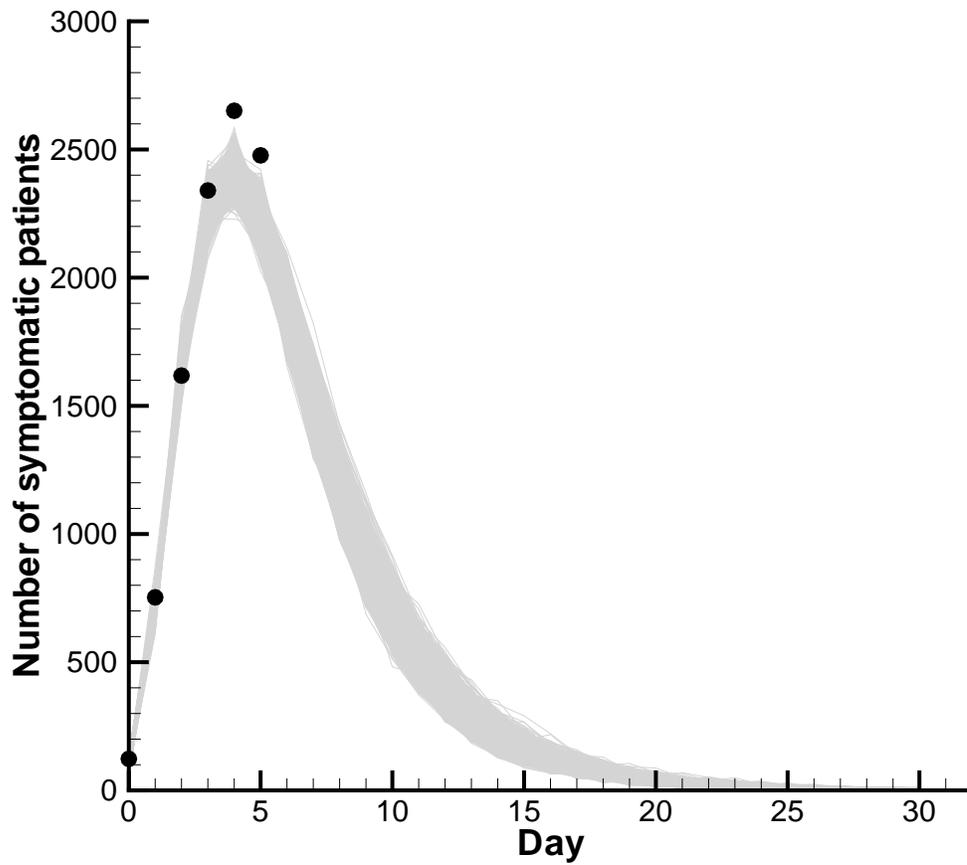


Figure 3.3. Evolution of the epidemic using 4000 samples taken from the joint PDF of the attack parameters based on the corrected data on Day $R = 6$ as presented in Figure 3.1. The actual number of people showing symptoms by Day 6 is plotted using symbols.

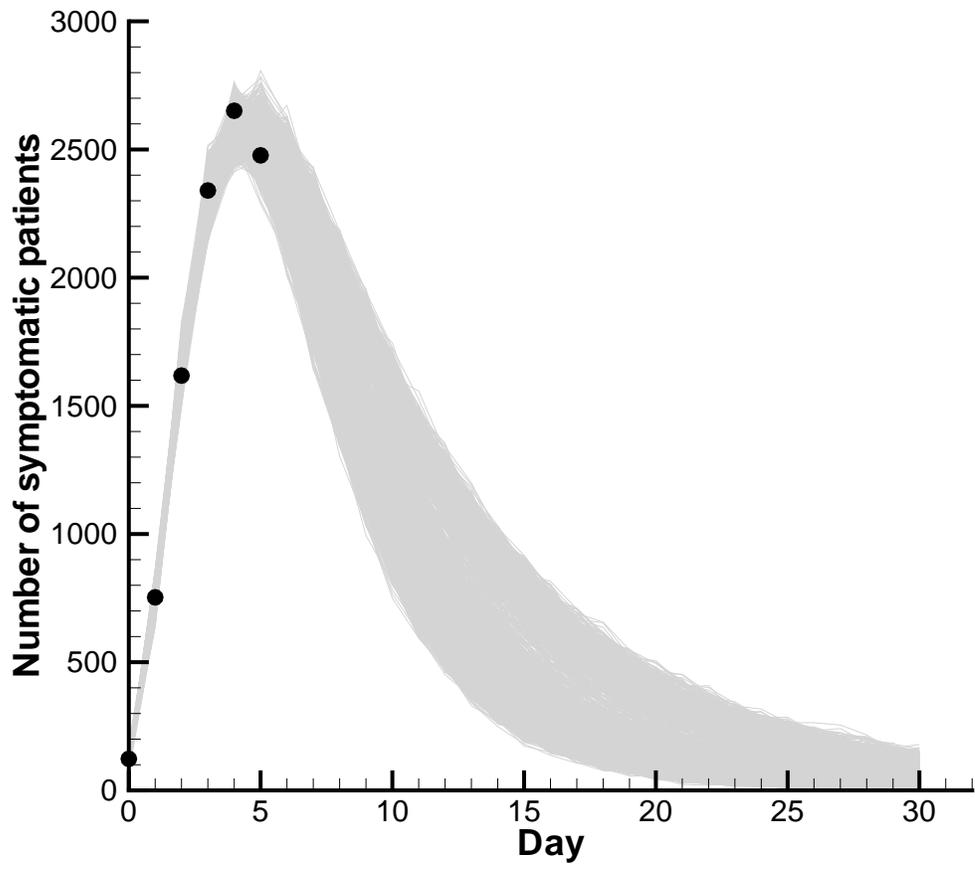


Figure 3.4. Evolution of the epidemic using 4000 samples taken from the joint PDF of the attack parameters based on the pristine data, for comparison to Figure 3.3. The actual number of people showing symptoms by Day 6 is plotted using symbols.

This page intentionally left blank.

Chapter 4

Resource allocation in attacks with non-communicable agents

In this section, we formulate a resource allocation problem, based on probabilistic reconstruction of an attack’s parameters, as discussed in [14]. The solution of the problem results in a time-dependent allocation profile, which can be updated with the availability of information and constrained by the actual capabilities of the transportation infrastructure. This enables an efficient yet realistic allocation of resources in a “reload” scenario. In this chapter, we will use anthrax as the agent for conducting the attacks.

4.1 Formulation

The problem of resource allocation in reload scenarios is governed by two non-dimensional numbers, which are ratios of timescales. The response to a bioattack, of a reload nature, is governed by three main processes viz,

1. the time-scale of the epidemic/outbreak, τ_E . For the anthrax attacks, $\tau_E \approx 15$ days.
2. the time-scale of the transportation infrastructure, τ_R . This may be the time required to gather and transport significant fractions of the resource demand, starting from a “standstill”. Typically $\tau_R \approx 2$ days.
3. the stagger time-scale τ_S . This is the average time-delay between subsequent attacks.

The following time-scale ratios should hold if a set of staggered bioattacks are to be amenable to the resource allocation techniques outlined here:

- $\tau_S/\tau_E \ll 1$. This condition indicates that the stagger should be relatively small, e.g., typically less than the time required for the first attack to be detected. It is expected that the heightened security posture after the *detection* of a bioattack will prevent any subsequent attacks. Also, the parameter domain $\tau_S/\tau_E \gg 1$ indicates a large separation between two bioattacks; these can be addressed separately and do not constitute a “reload” scenario. A “good” value for τ_S/τ_E is 0.2.

- $\tau_R/\tau_I \ll 1$. This condition indicates that response time of the transportation infrastructure should be significantly smaller than the epidemic, so as to accommodate changes in resource allocations, as dictated by the data stream. An inability to respond to the dynamic changes in resource demand essentially renders the current treatment useless; one may as well consider a static/point estimate of the resources required, allocate it, and ignore the information content of the morbidity streams. A “good” value is $\tau_R/\tau_I \approx 0.2$.

4.1.1 Optimal allocation of resources

As noted above, the basic problem is one of making optimal allocation of resources under significant uncertainty. In previous work, we showed how to capture this uncertainty in a PDF that can be used to predict the number of patients who arrive at the hospital requesting treatment. Specifically, we construct a number of such scenarios that are consistent with the data we currently have.

The basic strategy is to use the data that we have to make a decision of how much resource to allocate today. We also estimate the quantity of resources that might be required in the future (i.e., a resource allocation profile) to plan future logistical requirements. When new information is obtained the next day, a new estimation of the attack parameters (and the resource demands) is performed. The scenarios are recalculated in light of the new data, the available resources are decreased by the amount allocated today, and the allocation recalculated. Thus we have one decision variable: the amount we allocate today based on the information at hand. Let that variable be a .

We adopt the following assumptions and notations: Let

- K be the number of scenarios;
- T be the number of days that we consider, i.e., the planning horizon of the epidemic;
- $N_{k,j}$ be the number of people requiring treatment who arrive on day j in scenario k ;
- $r_{k,j}$ be the allocation made on day j in scenario k ;
- $D_{k,j}$ be the number of people who die on day j in scenario k ;
- D_k be the total number of people who die in scenario k ;
- R be the total number of resource units available for the attack, where we assume for simplicity that one unit of resource treats one patient;
- $s_{k,j}$ be the number of resource units available on day j of scenario k ;
- t_j be the fraction of people, arriving on day j , who will die having been treated;
- u_j be the fraction of people, arriving on day j , who will die not having been treated.

We assume that those arriving later in the attack will be more likely to be successfully treated. This is motivated by the fact that longer incubations generally indicate a lower dose exposure (or a robust constitution). Thus we assume that $t_{j+1} > t_j$ and $u_{j+1} > u_j$. In practice, we make this difference ($t_{j+1} - t_j = u_{j+1} - u_j = \varepsilon = 10^{-6}$) small and it merely serves as a mathematical stratagem to remove multiple solutions.

To construct the optimization problem, we need to specify the objective function. As a first cut, let us assume that we want to minimize some function of the sum of the number of deaths in each scenario, i.e., we seek to minimize

$$\sum_{k=1}^K M(D_k)$$

where $M(D_k)$ is some measure of D_k . We could consider various measures, but clearly one could take M to be simply the expected number of deaths. The optimization problem is then

$$\min_a \sum_{k=1}^K M(D_k),$$

where we have to specify constraints on resources and on how to compute D_k . The resource constraint is, clearly,

$$0 \leq a \leq R.$$

Given the treatment assumptions described above, we can easily compute

$$D_{k,j} = r_{k,j}t_j + (N_{k,j} - r_{k,j})u_j. \quad (4.1)$$

For day 1, we tentatively substitute a for $r_{k,1}$.

The allocations $r_{k,j}$, $j > 1$ can be chosen to be the optimal allocations for scenario k , given that allocation a was made in day 1. These allocations will be constrained as follows:

$$\begin{aligned} r_{k,j} &\geq 0 \\ \sum_{j=2}^T r_{k,j} &\leq R - a \text{ for each } k. \end{aligned} \quad (4.2)$$

Although it is possible to iteratively solve problems for each scenario separately, it is more efficient to make the collection $r_{k,j}$ variables in the optimization problem and solve one large problem rather than K smaller problems for each trial value of a .

Before we pose the final version of the initial problem, we must address an important situation. It is possible that in some scenarios, an allocation a or $r_{k,j}$ will be greater than the number of people who arrive, i.e., $a > N_{k,1}$, in which case, the value of $D_{k,j}$ from above will not correctly calculate the number of deaths. To handle this situation, instead of substituting a for all $r_{k,1}$, we retain separate $r_{k,1}$ variables and impose the constraint $r_{k,1} \leq a$ for all k . We also impose the simple bound constraints $r_{k,j} \leq N_{k,j}$ for all (k, j) , and we change (4.2) to

$$\sum_{j=2}^T r_{k,j} \leq R - r_{k,1}.$$

Another important concern is that without further constraints, the optimal choice of a may be to allocate all possible resources on the first day, which seems unlikely to be the best policy. One way to address this issue is to make tentative allocations for all days in the planning horizon, i.e., to introduce decision variables $a_i \geq 0$ for $1 \leq i \leq T$ and to restrict each scenario's allocations by $r_{k,j} \leq a_j$, with $a_1 = a$ and

$$\sum_{j=1}^T a_j \leq R.$$

In other words, we decide, *a priori*, that the daily allocation cannot exceed a certain level. Obviously the level chosen has a significant impact on the quality of the allocation calculated. This is studied further below.

Of course, the purpose of the exercise is still to choose the first day's allocation $a = a_1$. Another possibility is to penalize over-allocation of resources, in keeping with some standard approaches. To do this, we introduce a penalty term in the objective function of the form

$$\rho \cdot (a - r_{k,1})_+,$$

where $x_+ = x$ if $x > 0$ and 0 otherwise, and ρ is a constant chosen to appropriately balance the costs, i.e., penalize wastage / overallocation of resources. We choose this form for our studies here.

Policy makers may further wish to limit daily allocations to specified fractions of the available resources, say $a_j \leq \sigma_j R$. For simplicity, below we use a common value $\sigma_j = \sigma \in (0, 1]$ for all j (with $\sigma = 1$ imposing no further restriction).

The final topic we consider here is the reuse of resources. As noted, a high percentage of patients being treated will die anyway and they will die at a nonuniform rate. Some, in fact, will die quite early and their resources can be used on incoming patients. Data for estimating the rates are not readily available, but reasonable approximations can be made. Based on typical treatment progressions, the longer one survives, the more likely complete recovery becomes. Thus the percentage of

people who die after n days of treatment should increase rapidly for a few days and then gradually decrease. As a first cut, we assumed a ten-day period and used a simple function, $f_n = f_n^0 / \sum_{k=1}^{10} f_k^0$ with $f_n^0 = n / (1 + \exp(n/2))$ to estimate these rates (with $f_n = 0$ for $n > 10$). This is in the form of the expected percentage of people being treated who will die n days after treatment has begun. Given f_n , it is straightforward to estimate the number of resources that will be available on any given day, as in (4.5) below. Along the same lines, as noted above, there will be some scenarios for which allocations will exceed demand and the extra resources will likewise be available for incoming patients. The number of people who will die is still given by (4.1).

Our optimization problem is shown in Figure 4.1. Some remarks about it are in order. It is a two-stage stochastic optimization problem with recourse. The first stage is today and the second stage is days 2– T . Each scenario takes recourse on the basis of today’s allocation and does the best that it can after that. Constraints (4.3), (4.4), and (4.5) together imply that each scenario consumes at most R resources.

$$\begin{aligned}
 & \min_{a, r_{k,j}, s_{k,j}} \frac{1}{K} \sum_{k=1}^K \{M(D_k) + \rho(a - r_{k,1})_+\} \\
 & \text{subject to: } 0 \leq a \leq \sigma R \\
 & \quad 0 \leq r_{k,j} \leq \min(N_{k,j}, \sigma R) \\
 & \quad r_{k,1} \leq a \\
 & \quad r_{k,j} \leq s_{k,j} \quad (4.3) \\
 & \quad s_{k,1} = R \quad (4.4) \\
 & \quad s_{k,j} = s_{k,j-1} - r_{k,j-1} \quad (4.5) \\
 & \quad \quad \quad + \sum_{n=1}^{j-1} f_n t_{j-n} r_{k,j-n} \\
 & \quad D_{k,j} = t_j r_{k,j} + (N_{k,j} - r_{k,j}) u_j \\
 & \quad D_k = \sum_{j=1}^T D_{k,j}.
 \end{aligned}$$

Figure 4.1. Multi-scenario resource allocation problem.

One could, in principle, construct a multi-stage problem by dividing the days 2– T into two or more stages. Suppose, for example, that the second stage is days 2–4. Then one could trace each of the K scenarios through day 4. At that point, one assumes that, for each k , the data $N_{k,j}$, $j = 1, \dots, 4$, are “true”, constructs a PDF based on this data and samples that to obtain K new time series for each k . Although this can be easily continued, it is clear that the number of possible paths through the

attack grows rapidly. In this paper, we confine ourselves to just the two stages, but the extension to more stages is theoretically possible.

We have not yet specified M in the objective function, but we note that the constraints are all linear. Thus, if M is a linear function, we have a classical linear programming problem for which there are many excellent algorithms available. If we take M to be the identity operator, i.e.,

$$M(D_k) = D_k = \sum_{j=1}^T D_{k,j} \quad (4.6)$$

then we are simply computing the expected number of deaths in each scenario, and

$$\bar{D} = \frac{1}{K} \sum_{k=1}^K D_k \quad (4.7)$$

is the expected number of deaths over all the scenarios. This has an obvious appeal; results using this choice of M are reported in the next section. A potential problem with this is that scenarios with a large number of infected people could dominate the decisions. Recall our assumption that people arriving later are better candidates for treatment; in a scenario with a large number of people, the algorithm would delay the allocation of resources much more so than for a scenario with a much smaller number of infected people. It could be argued that the sampling procedure should properly account for this, but one could also divide D_k by the total number of people infected in scenario k . This downplays the influence of the larger cases, while keeping the problem linear.

A different approach, related to the work in [37], is to compute the optimal number of deaths for each scenario in K separate problems. Call the results D_k^* . Then one could obtain an allocation that stays as close as possible to all of these in some sense. A natural way to do this is to minimize the variance between the vector D_k^* and the vector D_k resulting from any other allocation. In particular, one could use

$$\sum_{k=1}^K (D_k - D_k^*)^2$$

as the objective function. In [37] we referred to this as the least-regret formulation with the interpretation that the allocation made today is the one that we will least regret in the future since it does reasonably well for all scenarios. As above, we could scale each of the terms by the total number of arrivals in that scenario. Since this is a quadratic function, the optimization is now a convex quadratic programming problem; again, good algorithms exist. The computation of each D_k^* is a small linear programming problem that is solved quickly. One advantage of the least-regret formulation is that it produces an allocation schedule, i.e., an allocation for each day over the entire course of the episode. This provides the emergency manager with a better planning aid than just the allocation for today. Thus, in our reporting below, we calculate the allocation for today using the problem in figure 4.1 and then calculate the schedule for the remaining days using least regret.

We now present some numerical results illustrating some of the issues raised here.

4.2 An attack on one city

To explore the approach, we first generated a test case involving an anthrax attack on a single city. This is described in [37]. Briefly, an aerosolized anthrax release is simulated over a domain with spatially variable population density. Per this distribution and an atmospheric dispersion model, 22,384 individuals are infected with a range of doses, with an average dose of 1470 spores. People develop symptoms over time; the time series for the first 10 days is $\{3, 123, 719, 2046, 2202, 2194, 2058, 1918, 1656\}$. This time series was used to draw 100 samples from the joint PDF of the attack parameters using a single-component random-walk Markov Chain Monte Carlo (MCMC) sampler. Note that these samples were drawn after the MCMC sampler had “burnt-in” and had “converged” per the `mcgibbsit` package in R (Chapters 7 and 8 in [38]; also see [39]). For each attack parameter sample, 10 epidemic realizations were calculated (the forward model is stochastic), resulting in a set of 1000 epidemic realizations (or scenarios). Such ensembles, generated from the first 5 days of data in the time-series above, are plotted as the gray region in Figure 4.2. Note that we measure time from the day that the first person was diagnosed with anthrax (rather than the time of attack/infection). The distribution developed with data collected through Day 7 is much narrower than that through Day 3, confirming that the addition of 4 extra days of data significantly reduces the uncertainty. This has not been plotted here.

The model was implemented in AMPL [40, 41, 42] and used the CPLEX 11 [43] simplex method to solve the problems.

We ran many tests based on the model described above. We fixed our available resources such that they could treat 10,000 patients (out of the 22,384 infected), i.e., they are scarce. Our first observation is that the form of the function M does not make much of a difference in the results. Thus all of the results we show here were calculated using (4.6) to minimize the expected deaths (4.7). Our second observation is that the penalty parameter, ρ , should be taken to be a small value to ensure its desired effect. After some tests with several values of ρ , summarized in Table 4.1, we chose use $\rho = 0.001$ for all of the results reported here.

Table 4.1. Effect of ρ on a and expected deaths (4.7).

ρ	a	\bar{D}
0.	10000	9358.0
0.0001	2384	9358.0
0.001	2364	9358.0
0.01	2317	9358.2
0.1	2261	9360.2

In Figure 4.2 we plot the allocations, given a resource demand drawn from 5 days of observations

in the time series. The gray region denotes the ensemble of scenarios. The time-series values used for the inference are plotted with triangles; the future observations in the time-series are plotted with diamonds. Allocations were calculated for $\sigma = 0.04$ and 0.1 . Clearly σ makes a significant difference. Recall that there are two possible reasons for imposing a constraint on the amount of resource that can be shipped on a given day: the first is that this may simply be a logistical constraint; the second is that the emergency manager may want to conserve resources as a hedge against a subsequent attack. Observe that our formulation only computes the allocation for Day 6; to give managers an idea of allocations that might be appropriate on subsequent days, we obtain tentative allocations for days 7– T by averaging the allocations for each day over all of the scenarios. (Subsequently arriving data should influence the actual allocations for later days.) As is evident in Figure 4.2, the severe restriction imposed by $\sigma = .04$ implies that many fewer resources can be allocated than for the lighter restriction of $\sigma = .10$. Thus there is a commensurate increase in the number of deaths with $\sigma = .04$, as we show in Figure 4.3. Here we plot the PDF of excess casualties (over the optimal/minimal level that we would have achieved had we perfect knowledge of the epidemic) for the two values of σ . As might be expected, the effect of σ (i.e., the placing of a ceiling on how much can be shipped on a given day) is felt mainly in those scenarios that project a large number of infected people turning symptomatic. We also see that increasing σ narrows the PDF (we reduce the long-tail probability of an extremely adverse outcome) while raising the peak of the PDF and moving to lower values of excess casualties, i.e., increasing the probability of a less adverse outcome. Since the probability mass under the PDF is 1, this is tantamount to increasing the probability of a certain (acceptable) level of casualties while simultaneously trading it to reduce the probability of an extremely adverse outcome — a classic hedging / risk management operation. This is captured quantitatively in the change of shape of the PDF with σ .

For $\sigma \geq 0.20$ we obtain an allocation (not shown here) resulting in very few excess casualties in each scenario. These results show that the model can be used for assessing the effects of conserving resources in anticipation of a second attack or for planning purposes to see the need for a higher shipping capacity.

One could also compare the PDFs of excess casualties if a “naive” approach to resource allocation was considered, e.g., given an R , one allocates on a scenario-by-scenario basis (leading to 1,000 allocations), then simply uses the mean of these allocations. Such a “naive” allocation results in a very long tail (see [37] for a comparison) and is not very competitive for hedging purposes vis-à-vis the more sophisticated techniques considered here and in [37]. For the rest of this paper, the “naive” approach will be ignored.

4.3 An attack on two cities: The “reload” case and the equilibration of pain

The main complication in dealing with an attack on two or more cities is in deciding how to allocate the resources among all of the cities. From the point of view of the model, it does not make any difference if a life is saved in the first city or the second. Thus, without further constraints, there is

an inherent non-uniqueness in the solution of the problem, since the optimal number of deaths can be achieved in many ways, including the extreme one of sending all of the resources to one city and ignoring the other. In practice, it seems reasonable to assume that there will have to be some “social” or infrastructural constraint to ensure that all cities are treated fairly. We illustrate how this could be achieved below, but first we deal with another issue, namely that of whether or not to anticipate subsequent attacks.

As noted in Section 4.1.1 an emergency manager may wish to restrict the amount of resources that can be shipped on each day. This is done by imposing the constraints $r_{k,j} \leq \sigma R$. The manager could equally well choose σ to conserve some of the resources in case there is a subsequent attack, the “reload” case. If there is a subsequent attack, there is no way to say anything about it until there is some evidence in the form of people in the second city arriving at the local hospital in need of treatment. As is the case for the first city, a few days of data are required before any reasonable PDF can be computed and sampled.

Extending the basic model above to the case of several cities is straightforward. The major addition for the reload case is the social constraint. We illustrate the possibilities with a simple constraint that seeks to ensure that each city receives a proportional amount of the resources. A way to do this is to impose the constraints

$$D^i/A_i \leq (1 + \pi) \sum_{j \neq i} D^j / \sum_{j \neq i} A_j,$$

where D^i are the deaths in city i , A_i is the total number of patients in city i , and $\pi \in [0, 1]$. For the results reported here, we used $\pi = 0.1$, so that the relative resource allocations are within 10%.

We demonstrate this allocation approach on a simulated reload scenario. The first attack (on City A) is the same as in Section 4.2. However, on Day 3 of the first attack, City B records an anthrax diagnosis and it is verified that it too has been attacked. The time-series for City B is $\{0, 0, 1, 76, 711, 1765, 2720, 3099, 3186, 2896\}$ for the first 10 days. The attack on City B was simulated in the manner described in [37]. 29,861 people were infected, with an average dose of 2749 spores. The two attacked cities therefore have a resource demand of around 50,000 units. In the study below, we will assume that only 25,000 units are available.

The allocations are shown in Figures 4.4 and 4.5 for Day 6 of the attack, i.e., we have a time-series 5 days long for City A and 3 days long for City B. The gray region in Figures 4.4 and 4.5 show the ensemble of scenarios for the two attacks; as expected, the ensemble for City B is far broader than City A, denoting a larger uncertainty arising from a smaller time-series of observations. The observed and unobserved evolution of the epidemic in the two cities is plotted using triangles and diamonds. The allocations developed with $\sigma = 0.04$ and 0.1 are plotted for Day 6 (and beyond) of the epidemic. Note that the allocation is only meant for Day 6. Both the plots demonstrate how allocations are curtailed as σ decreases, leading to extra casualties, especially for scenarios that project larger epidemics. Also note that the effect of σ is felt mostly during the peak of the epidemic; the allocations are similar towards the end. This is a consequence of our modeling decision to slightly favor later allocations.

In Figure 4.6 we plot the PDFs of excess casualties (over the minimum that we would achieved had we perfect knowledge of the attack and the epidemic). The excess casualties for Cities A and B, for $\sigma = 0.02, 0.04$ and 0.1 are totaled and plotted. Note that the $\sigma = 0.1$ case is not at all restrictive and one even has overallocation of resources (the “negative” casualties). This happens when two exceptionally small scenarios for City A and B are considered. Note that the σ value merely places a bound on daily allocation; the constraint that daily allocations must add up to the available resources is not violated. The hedging effect of σ seen in Section 4.2 is also reproduced here, though with a few modifications. In all cases, we see a multimodal excess-casualty distribution. While $\sigma = 0.04$ does manage to translate the excess-casualty PDF to the left (vis-à-vis $\sigma = 0.02$), we see the width of its support is unchanged, i.e., the higher value of σ reduces the expected casualties (and consequently risk), but does not improve the hedge compared to $\sigma = 0.02$.

4.4 An attack on one city with corrected data

As reported in Chp. 3, the raw data consisting of just the number of people who arrive on a given day is not sufficient to create good PDFs, since we need to know the day on which they became symptomatic. In this section, we show the differences in the PDFs between the corrected and uncorrected data and then show a series of resource allocations made based on the corrected data over a 6-day period of the same attack. The results demonstrate that the resource allocations early in the attack are quite good. That is, even as we get additional data and the PDFs narrow, the allocations do not differ very significantly.

Fig. 3.3 and 3.4 show the resulting set of scenarios from Day 6 with the corrected and uncorrected data. As one can readily see, there is a significant difference in the size and range of the scenarios and thus there would be a significant difference in the allocations.

In Fig. 4.7 we show the allocations schedules that are calculated for days 4–9. We assume that it takes 3 days to ramp up the transportation infrastructure for resource distribution purposes, thus placing a constraint on the allocation that can be realistically performed. Days 4, 5, and 6 show the effect of this ramp-up to enable a full allocation of 10% of the resources. Note that the form of these allocations is nearly the same, i.e., we ship at high levels early in the attack and then decrease the shipments rapidly thereafter. Also note that the scenarios for days 8 and 9 are very close and so we do not expect much change to occur after day 9. Also note that in this case, a better reconstruction of the epidemic (with more data) results in lower levels of allocation (see the allocations for Day 6-9). Further, most of the allocation is done early with a quick curtailing of allocation later in the outbreak. This is because as the outbreak is better defined/reconstructed from data, the fall-off in the epidemic curve becomes more certain, allowing a better (less uncertain) allocation in that regime.

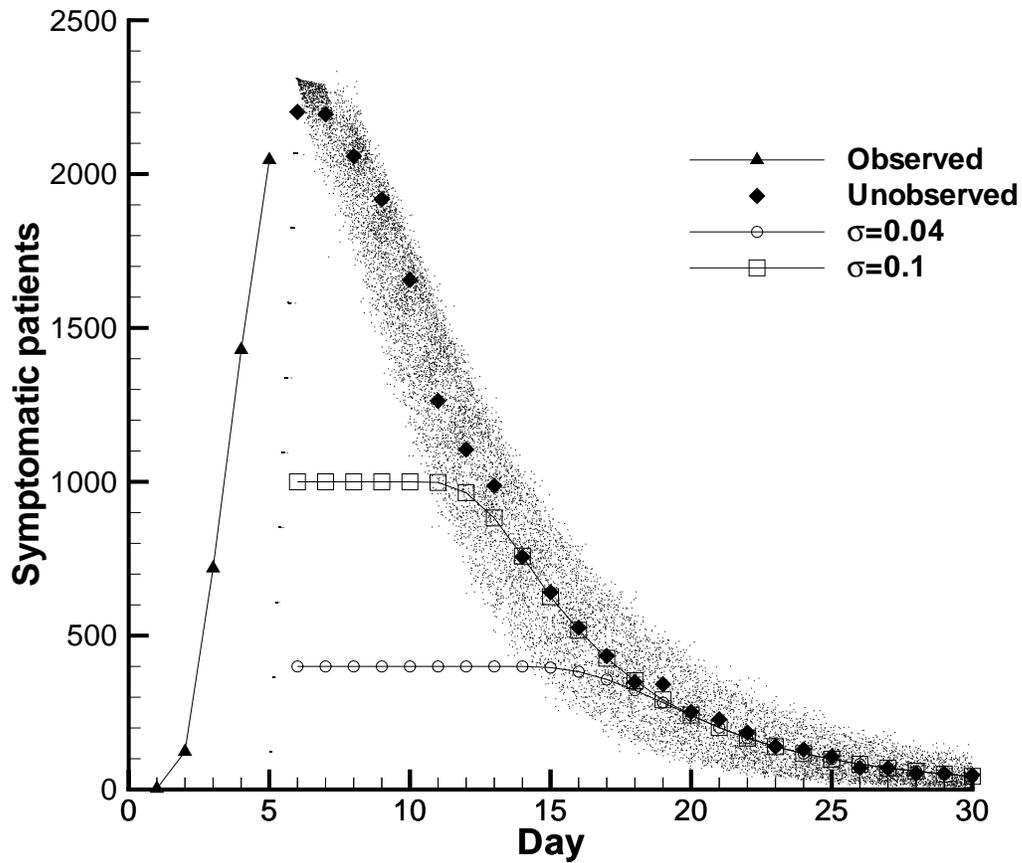


Figure 4.2. Allocations for the attacked city, obtained from data collected over the first 5 days. The gray region denotes the evolution of all the scenarios considered. The net effect of σ is to reduce the allocation during the early days of the epidemic. $R = 10,000$. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds.

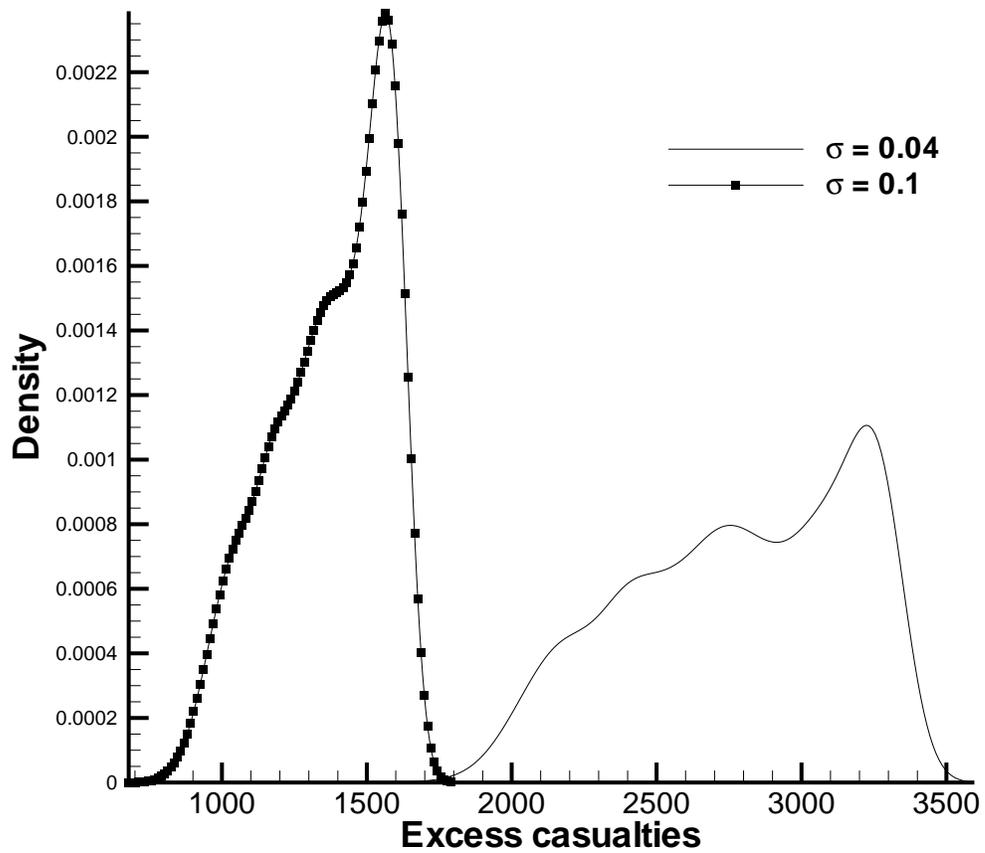


Figure 4.3. PDFs of excess casualties for $\sigma = 0.04$ and 0.1 . $R = 10,000$. A tighter daily constraint on allocations ($\sigma = 0.04$) increases the probability of excess casualties. However, note that the PDFs have rather compact support.

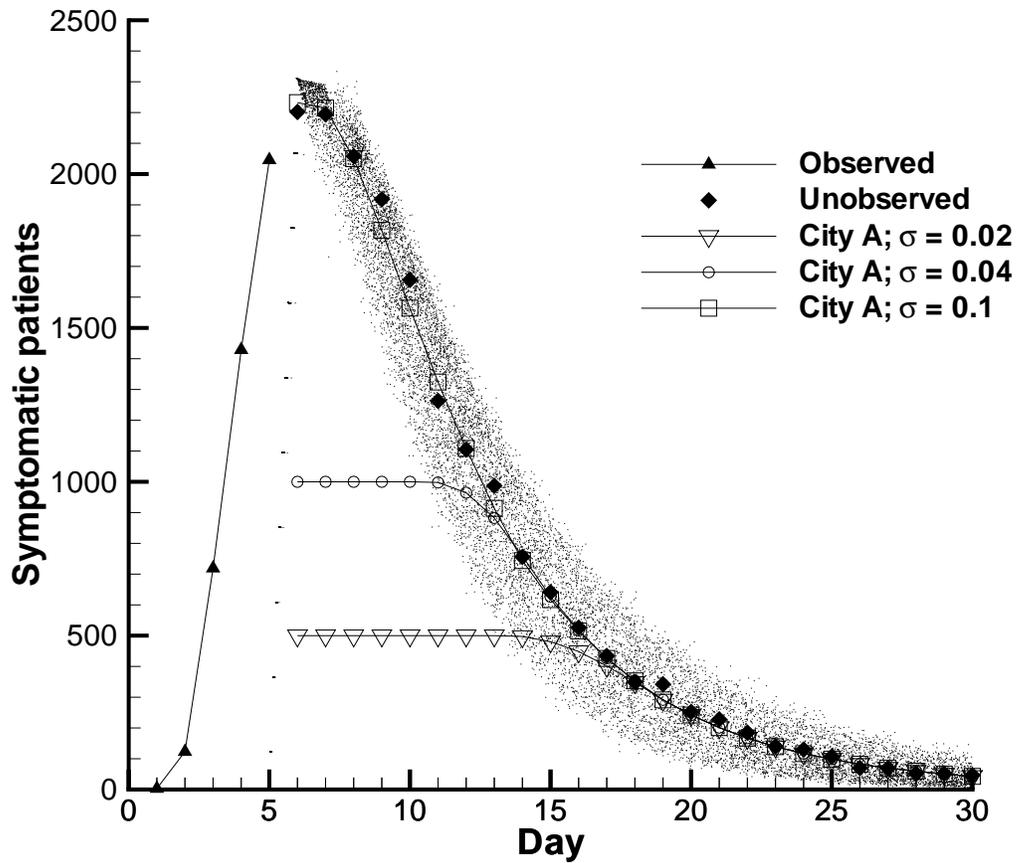


Figure 4.4. Allocation under various values of σ for City A. The gray region denotes the evolution of all the scenarios considered. $R = 50,000$ (total for both cities). These allocations were drawn from data collected over 5 days; allocations are for Day 6. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds.

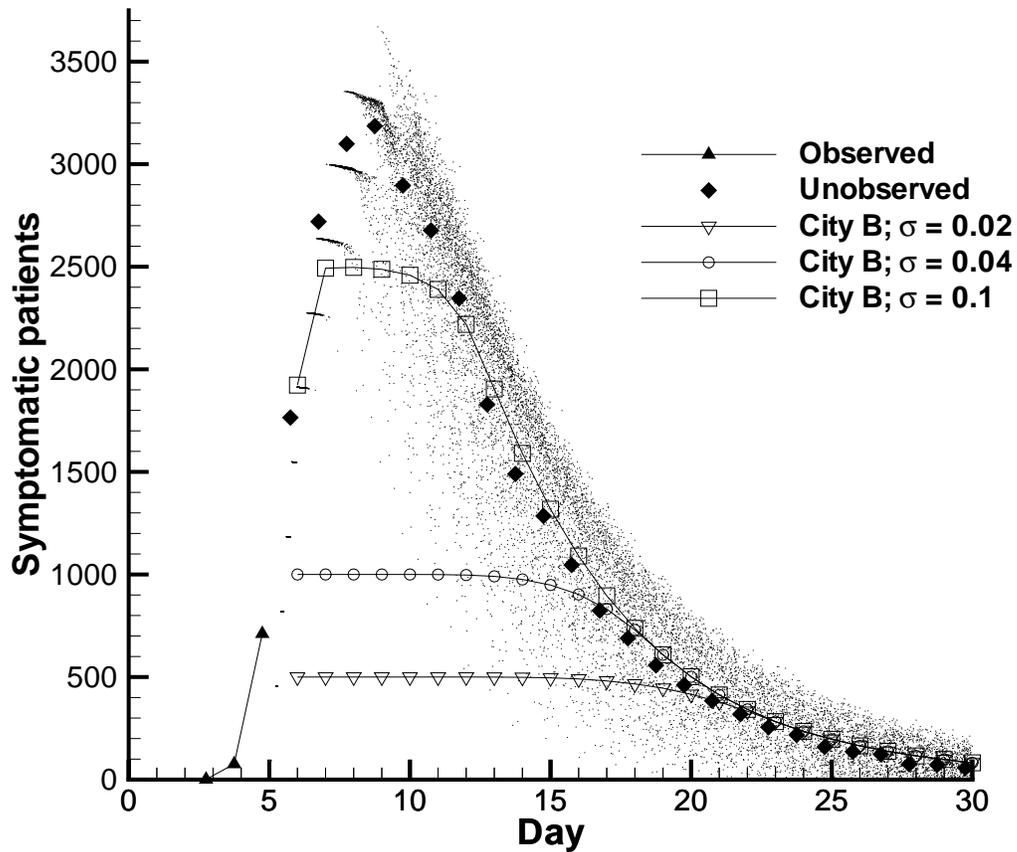


Figure 4.5. Allocation under various values of σ for City B. $R = 50,000$ (total for both cities). These allocations were drawn from data collected over 3 days (Days 3, 4 and 5, the attack stagger being 2 days); allocations are for Day 6. The observed evolution of the epidemic is plotted with triangles; the future, unobserved evolution with diamonds. The gray region denotes the evolution of all the scenarios considered.

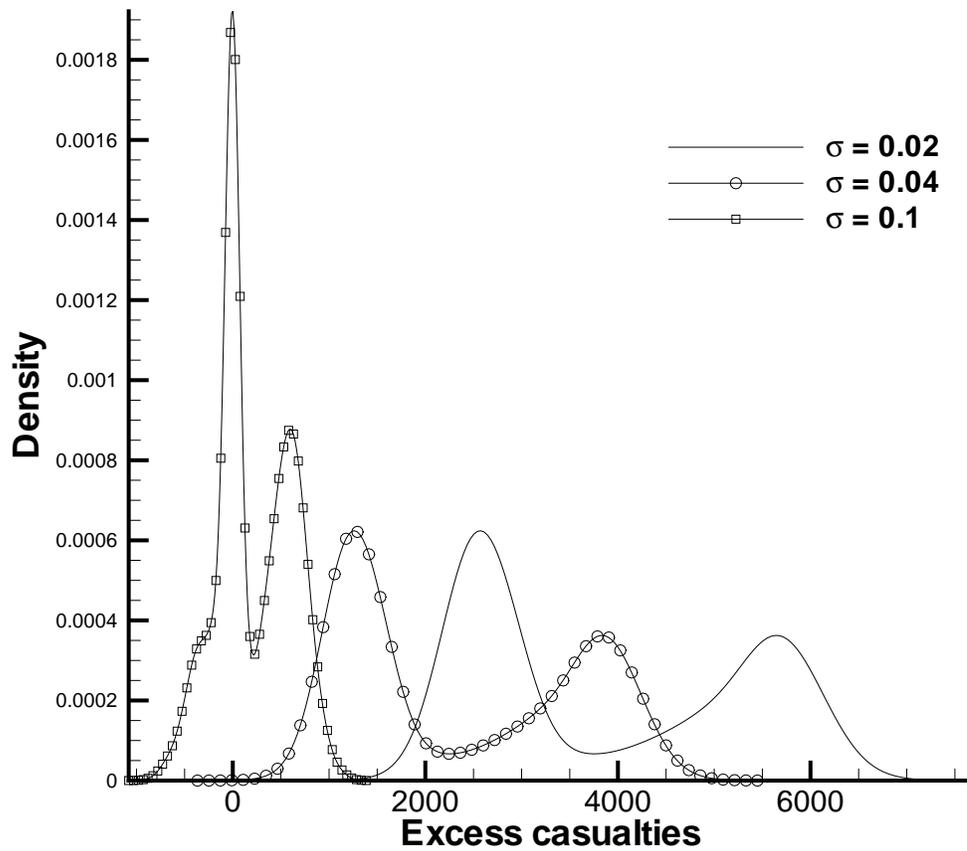


Figure 4.6. PDF of excess casualties for $\sigma = 0.02, 0.04$ and 0.1 . $R = 50,000$. Note how $\sigma = 0.1$ results in certain “negative” casualties, i.e., excess resources in certain low casualty scenarios. Most of the excesses occur in the smaller attack (City A). The PDF corresponding to $\sigma = 0.02$ and 0.04 have a similar support widths but the higher value of σ reduced the expected value of casualties.

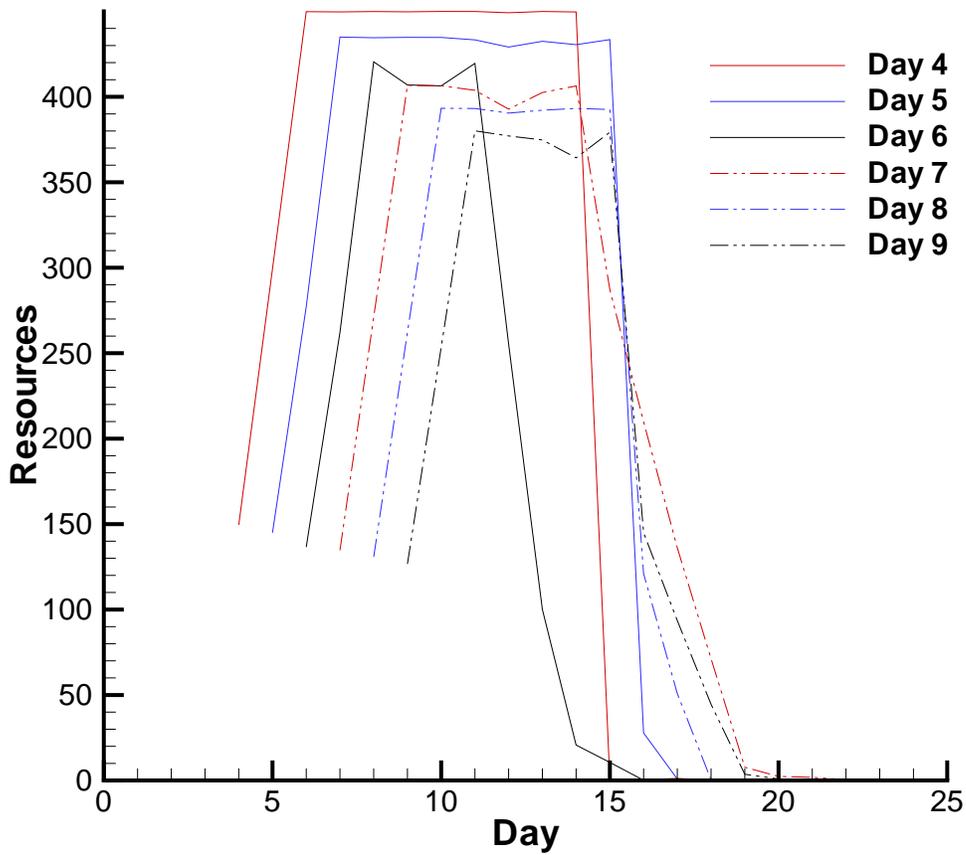


Figure 4.7. Resource allocation profiles over the durations Days 4-9 calculated from inferences drawn from increasing amounts of data. We see the effect of the 3-day ramp up. Further, most of the allocation is done early with a quick curtailing of allocation later in the outbreak.

Chapter 5

Resource allocation for outbreaks caused by communicable diseases

In this chapter, we will address the problem of resource allocation in case of outbreaks of communicable diseases. As mentioned before, the process consists of two elements viz. estimation of outbreak parameters and the resource allocation process under an uncertain characterization of the outbreak (in the form of an ensemble of outbreak realizations).

The resource allocation problem for communicable diseases is considerably more involved compared to non-communicable diseases. In Chp. 4, where we discussed the allocation of equipment and personnel in response to an anthrax attack, we assumed that the level of resource allocated made no difference to the evolution of the outbreak. This is because, unlike antibiotics which suppress and can cure the disease, the resources considered (equipment and personnel) provide palliative or supporting care. Thus the resource allocation / optimization procedure did not require estimating the impact of the allocation on the outbreak itself; only the effect on casualties was estimated. If any distribution of antibiotics disrupted the outbreak, the data-driven methodology would capture the disruption and predict a smaller outbreak (though with a time-lag, required to collect sufficient data to capture the disruption).

In case of a communicable disease, equipment and personnel are most likely to be engaged in disrupting the transmission itself, and any resource allocation procedure will require the evaluation of its disruptive effect. This leads to an extremely computationally intensive procedure, which we will not address in this study. However, there also exist resources, e.g., drugs like Tamiflu, which provide palliative/supporting care, whose demand estimates are set by the size of the outbreak; the availability of such resources reduces casualties (and other losses) but does nothing to disrupt the epidemic. However, if exogenous process (e.g., medical responses like vaccinations etc) are involved in modulating disease spread, their effect can impact resource allocation greatly. Thus such modulations need to be captured for a proper estimation of resource demand.

In our study we will concentrate on resources that provide palliative care during an epidemic of a communicable disease. The procedure developed for anthrax in Chp. 4 is directly applicable, and we will not demonstrate its performance such epidemics. Rather, we will develop an estimation procedure for the important epidemic parameters, with the understanding that posterior predictive calculations (which result in an ensemble of outbreak realizations) and the least-regret calculations are identical to that of anthrax and therefore need not be investigated further.

5.1 Disease dynamics

Epidemics of communicable diseases show rich dynamics. Starting from a few index cases, they spread. In certain cases, the outbreak may cease because of its low transmissivity or because it infects a relatively unconnected clique of people. In other cases, it spreads widely, till changes in social behavior (usually social distancing) halts its spreads. Thus the infection intensity (rate of new infections) initially increases in time, before settling into a decline. This temporal variation of infection intensity cannot, of course, be observed; at most one may know the number of people exhibiting symptoms, at the end of their incubation period. Typically (e.g., for smallpox, plague, influenza and a host of diseases), the incubation phase is not contagious, and transmission starts only after a person shows symptoms. The symptomatic are often measured/recorded when they seek medical care; if this data is used to analyze an outbreak, one must also accommodate a reporting delay, as described in Chp. 3.

In this chapter, we will attempt to infer outbreak parameters of a communicable disease from a time-series of symptomatics (and NOT people seeking medical care, i.e., we will not include reporting delay in our inference) collected on a daily basis. Thus we have a time-series n_i of the number of new symptomatics on Day i , over a time duration $0 \leq t \leq T$. We assume that there exist a total of N_{tot} affected people, of which a fraction $1 - \alpha$ are index cases. The index cases are assumed to have been infected τ days before the first exhibition of symptoms, i.e., $\tau < 0$. There exists an unknown infection intensity $f(t; \mathbf{p})$, parameterized by \mathbf{p} . The objective is to determine estimates of $\{N_{tot}, \alpha, \tau, \mathbf{p}\}$. Choosing the form of $f(t; \mathbf{p})$ is a challenge; sometimes the choice may change as an epidemic progresses.

The data for the inference will be obtained from agent-based simulation of a communicable disease. The technique depends on the existence of a social network between individuals (agents), over which the spread of the disease occurs. A outbreak is very dependent on the index cases (or rather the connectivity of the index cases in the social network) and multiple simulations with the same *number* of index cases (but with different choices of them) can lead to very different outbreaks. The simulation technique is very similar to the one employed in [44] and is described in detail in [23], Sec. 3.

Below, we formulate a Bayesian inverse problem, where we leave $f(t; \mathbf{p})$ unspecified. Thereafter, we demonstrate the inference technique on a plague and an influenza epidemic. In each case, we describe the particulars of the transmission dynamics of the outbreaks. The true values of N_{tot} , τ and $1 - \alpha$ are known from the simulation and are provided for comparison with inferred values. The estimates of \mathbf{p} as a function of n_i are also provided.

5.2 Formulation of the inverse problem

Consider an epidemic that has been observed in the time duration $0 \leq t \leq T$, during which time, $N_{tot} = N_{ind} + N_{sec}$ have been infected. This includes the N_{ind} index cases and N_{sec} secondary cases. For large T , $\alpha = N_{sec}/N_{tot} \approx 1$. Assume that the index cases were infected at time $\tau, \tau < 0$. $t = 0$

indicates the time the first person (one of the index cases) shows symptoms and starts spreading the disease.

During the time-period $[0, T]$, there exists a time-dependent infection intensity given by $\alpha N_{tot} f(t; \mathbf{p})$, which is the rate at which people are infected. The function $f(t; \mathbf{p})$ is unknown and models the spread of the disease, which in turn is governed mainly by the social network and the transmissivity of the disease. Note that

$$\int_0^T f(t; \mathbf{p}) dt = 1 \quad (5.1)$$

The limits of integration are $[0 \dots T]$ since this is the time-duration over which symptomatic patients spread the disease. Consider the time-series $n_i, i = 0 \dots m$ of new symptomatic cases that appear in the time-interval $[t_{i-1}, t_i), t_i - t_{i-1} = \Delta t = 1$ day. n_i is a mixture of index cases and secondary cases turning symptomatic and can be given by

$$\begin{aligned} n_i^{model} &= N_{tot} \left((1 - \alpha) [C(t_i) - C(t_{i-1})] + \alpha \int_0^T f(s; \mathbf{p}) [C(t_i - s) - C(t_{i-1} - s)] ds \right) \\ n_i &= n_i^{model} + \varepsilon \end{aligned} \quad (5.2)$$

where $C(t)$ is the cumulative distribution function (CDF) of the incubation period of the disease and $\varepsilon \sim N(0, \sigma^2)$ is a measurement error.

Thus given a set of outbreak parameters $\{N_{tot}, \alpha, \tau, \mathbf{p}\}$, the likelihood of observing the time-series $n_i, i = 0 \dots m$ is

$$\pi(n_i, i = 0 \dots m | N_{tot}, \alpha, \tau, \mathbf{p}) = \exp \left(-\frac{\sum_{i=0}^m (n_i - n_i^{model})^2}{2\sigma^2} \right). \quad (5.3)$$

Using Bayes' theorem, the joint posterior probability $\pi(N_{tot}, \alpha, \tau, \mathbf{p})$ conditioned on data is

$$\pi(N_{tot}, \alpha, \tau, \mathbf{p} | n_i, i = 0 \dots m) = \exp \left(-\frac{\sum_{i=0}^m (n_i - n_i^{model})^2}{2\sigma^2} \right) \pi_{prior}(N_{tot}, \alpha, \tau, \mathbf{p}) \quad (5.4)$$

The posterior distribution can be sampled using a Markov Chain Monte Carlo (MCMC) method and marginalized to obtain probability density functions for each of the parameters in question. We use a simple random-walk MCMC to sample the posterior. Further, to assist in sampling we reparameterize the problem in terms of the logarithms of the parameters (except for τ which, since it is a negative quantity, is reparameterized in terms of $\log(-\tau)$). The priors on each of the parameters are assumed independent, are vague and are modeled as normal distribution, unless mentioned otherwise.

5.3 Test case 1: A plague outbreak

In this section, we consider the inference of a plague outbreak. First, we describe the epidemic model.

5.3.1 The outbreak simulation

The epidemic is assumed to evolve over a graph. Nodes in the graph represent people and the edges represent social links over which the disease could potentially travel. The transmission is stochastic and is modeled as a Poisson process with rate λ . Each node undergoes a susceptible – exposed (i.e., incubating) – infectious – removed sequence, with removal denoting recovery (and immunity) or death. The mortality rate for PPP, if left untreated, is 100% [45]. Treatment during the incubation phase has a 100% probability of success [45]. Treatment during the infectious (symptomatic) phase is unknown.

Each node i , on being infected, resides in the exposed and infectious phases for time τ_E and τ_I . τ_E and τ_I are random variables obeying a log-normal distribution with means (SD) of 4.3 (1.8) and 2.5 (1.2) days. These are obtained from [45].

The transmission model on a network is somewhat different from typical ODE-based SEIR models. In [46], it was observed by following infection networks that the effective reproductive number of PPP $R(T)$ could be expressed as

$$R(T) = R_0 \exp(-\delta T) \tag{5.5}$$

where $R_0 = 2.99$, $\delta = 0.0615$ and T is the time measured since the start of the epidemic. On the other hand, in [45], a “steady-state” reproductive number was assumed, and was found to be equal to 1.3. Curiously, if one averages the $R(t)$ over 30 days (the duration of the Madagascar and Mukden outbreaks considered in [46]), one obtains an average $\bar{R} = 1.39$. However, a reproductive number over 1.0 as proposed in [45] would indicate an epidemic that grows without bounds; on the other hand the expression in Eqn. 5.5 ensures that the epidemic will eventually die down.

Eqn. 5.5 is adapted for use in a network model. We proceed as follows. Consider a node i with incubation and infectious periods of τ_E and τ_I . Consider, too, that social links l_{ij} exist between nodes i and j , $j \in L_i$, where L_i is the set of nodes i is connected to (i.e., node i 's neighbors). Let $|L_i|$ denote the number of neighbors node i has.

When node i is infected, it is allocated a reproductive number per Eqn. 5.5. $R(T_i)$ denotes the number of people i will infect over the period τ_I , where T_i is the time that i was infected (this also ensures that the time-varying nature of the effective reproductive number is captured).

Since i has $|L_i|$ neighbors, a subset of them are marked for potential infection by i . We iterate through them and mark them for infection (via transmission from i) with probability $R(T_i)/|L_i|$. If L_i contains nodes which have already been infected, they are skipped over in the iteration process.

Once the potential “victims” of i are marked, we proceed with the dynamics of transmission. Consider a node $j, j \in L_i$ which has been marked for infection. The probability p_{ij} that i will infect j is given by $p_{ij} = 1 - \exp(-\alpha_i \lambda_i t_{ij})$ where $\lambda_i = 1/\tau_i$ is the rate of infection, t_{ij} is the time duration over which transmission could have occurred between i and j (essentially, duration of contact between i and j after i turned infectious) and α_i is a constant (for a given node i) that ensures that i will succeed in infecting j with a probability of 0.9999 (i.e., $1.0 - 10^{-5}$) by the conclusion of its infectious period. Thus $\alpha_i = -\ln(10^{-5})$.

5.3.2 Inference of outbreak parameters

Fig. 5.1 shows the temporal evolution of an outbreak. We start with 100 index cases, who infected an 1063 people over the next 43 days. Plotted in red are the new symptomatic cases, collated on a daily basis. It is noisy, and peaks at around day 15; thereafter, it decays, indicating a weakening epidemic. Plotted in blue is the latent infection intensity as a function of time. We see clearly that for the first 10 days, the infection intensity rises, after which it settles into a decline over the next 30 days.

Inferring the infection intensity $f(t; \mathbf{p})$ is key to predicting the evolution of the outbreak and consequently the resource requirements. We model the infection intensity as a Γ distribution, i.e.,

$$f(t; \mathbf{p}) = \frac{g(t; k, \theta)}{G(T; k, \theta)} = \frac{1}{\theta \gamma(k, T/\theta)} \left(\frac{t}{\theta}\right)^{k-1} \exp\left(-\frac{t}{\theta}\right). \quad (5.6)$$

where $\gamma(k, T/\theta)$ is the incomplete Gamma function, $g(\cdot)$ is the Gamma probability density function and $G(\cdot)$ is the corresponding CDF. Note that the expression in Eq. 5.6 obeys the normalization Eq. 5.1. Also, the parameters \mathbf{p} are the shape (k) and scale (θ) parameters of the Gamma distribution.

Following the description in Sec. 5.2, we perform an inference of the outbreak parameters using a time-series 20 days long. Over this duration, an extra 805 people were infected via transmission. The inference was performed using the transformed variables (i.e., the log-variables), though all results will be shown in terms of the variables $\{N_{tot}, \alpha, \tau, \theta, k\}$. The priors used are

$$\begin{aligned} \log(N_{tot}) &\sim N(\log(10^3), 10), \\ \alpha &\sim \mathcal{B}(1.25, 1.25), \\ \log(-\tau) &\sim N(0, 1), \\ \log(\theta) &\sim N(0, 2^2), \\ \log(k) &\sim N(0, 1) \end{aligned}$$

The first two prior are recognizable vague. The incubation period distribution for plague [45], with a median of 4.3 days, ensures that given 100 index cases, one may have, with high probability,

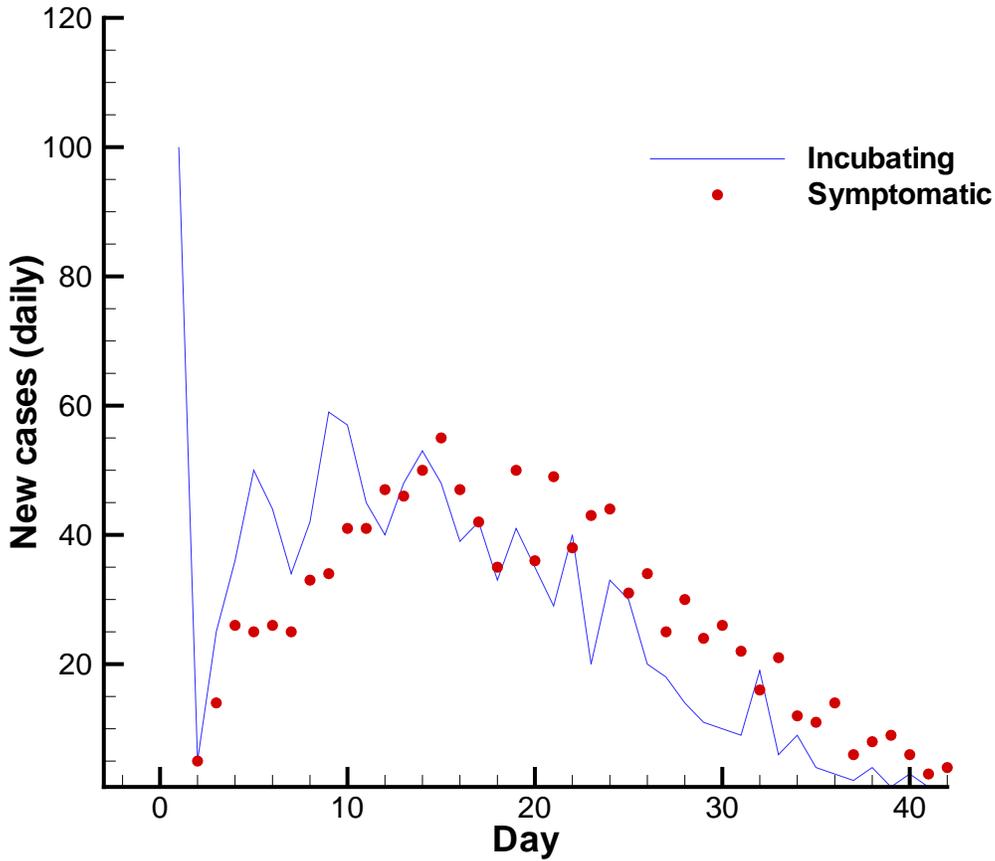


Figure 5.1. Evolution of the plague epidemic. New cases of symptomatic patients are plotted in red while the latent infection rate (daily infections) are in blue. The simulation was started with 100 index cases distributed in the social network. We see, from the symptomatic cases, that the epidemic peaks around Day 15. The infection rate peaks a few days before.

one symptomatic case within e days of infection. Similarly, the prior for θ , indicates that the time-scale for the decline of the epidemic may vary by e^2 days, approximately a week. The prior for k was chosen so that the rise in infection intensity would be roughly linear, a results that can be obtained from early-epoch linearization of conventional SEIR epidemiology models. The standard deviation of the measurement error, σ is set to 10.

In Fig. 5.2 we plot the MCMC chain and the histograms of the samples of the outbreak parameters, i.e., of $\{N_{tot}, \alpha, \tau, k\}$. The chains mix properly, i.e., ergodicity of the MCMC chain is achieved. The histograms on the right indicate the posterior marginalized distribution of the outbreak parameters. In Fig. 5.3, we determine the maximum *a posteriori* estimates of the outbreak parameters and plot

Table 5.1. Estimates (medians) of the outbreak parameters developed from time-series of different lengths. The 95% confidence intervals are mentioned in parenthesis and the true values are in square brackets.

Variable	m = 10	m = 15	m = 20	m = 25
N_{tot}	397 (284,609) [497]	673 (537, 1133) [725]	848 (743,974) [905]	1024.4 (927.6,1612.2) [1048]
α	0.79 (0.63, 0.95) [0.80]	0.87 (0.39, 0.96) [0.86]	0.90 (0.67,0.98) [0.89]	0.92 (0.41, 0.98) [0.90]
τ	-1.75 (-4.7, -0.4) [-2]	-1.9 (-10.6, -0.5) [-2]	-1.9 (-9.1, -0.5) [-2]	-2.0 (-13.8, -0.5) [-2]
θ	6.5 (0.79, 125)	10.4 (2.16, 156.6)	10 (3.5, 75)	10.2 (4.8, 36.5)
k	2.3 (1.2, 7.2)	2.13 (1.32, 4.7)	2.0 (1.25, 3.84)	1.98 (1.28, 4.98)

the corresponding latent infection intensity; the actual intensity too is plotted as a comparison. The blue dot at the left extreme is the number of index cases; the numerical estimate from the time-series has not been plotted. For both 10 and 25 days of data, we see that the infection intensity is properly captured, including the downturn in the infection intensity, as the outbreak begins to decline. The infection intensity curve from the 10-day time-series is seen to underpredict future infection intensities; the one developed from the 25-day time-series over-predicts it. This is a consequence of the Gamma-distributed model trying to capture the stochastic nonlinear dynamics of the spread of the disease on a social network. A more mechanistic model of the disease spread would likely provide a better fit, but it is unclear what such a model would be, which preserves the speed and parsimony of a Gamma model. In Table 5.1, we summarize the outbreak parameter estimates drawn from time-series of different lengths. Note that the true values of N_{tot} and α are functions of time, i.e., they increase as the epidemic progresses. We see that the inference procedure is fairly accurate; further, the simple model for the infection intensity is quite successful in summarizing the involved stochastic transmission dynamics simulated in the agent-based simulation.

5.4 Test case 2: An influenza outbreak

In Sec. 5.3, we showed how an outbreak that “failed to take off” could be modeled and inferred from a short time-series. The Gamma-distribution model of the infection intensity can be used to

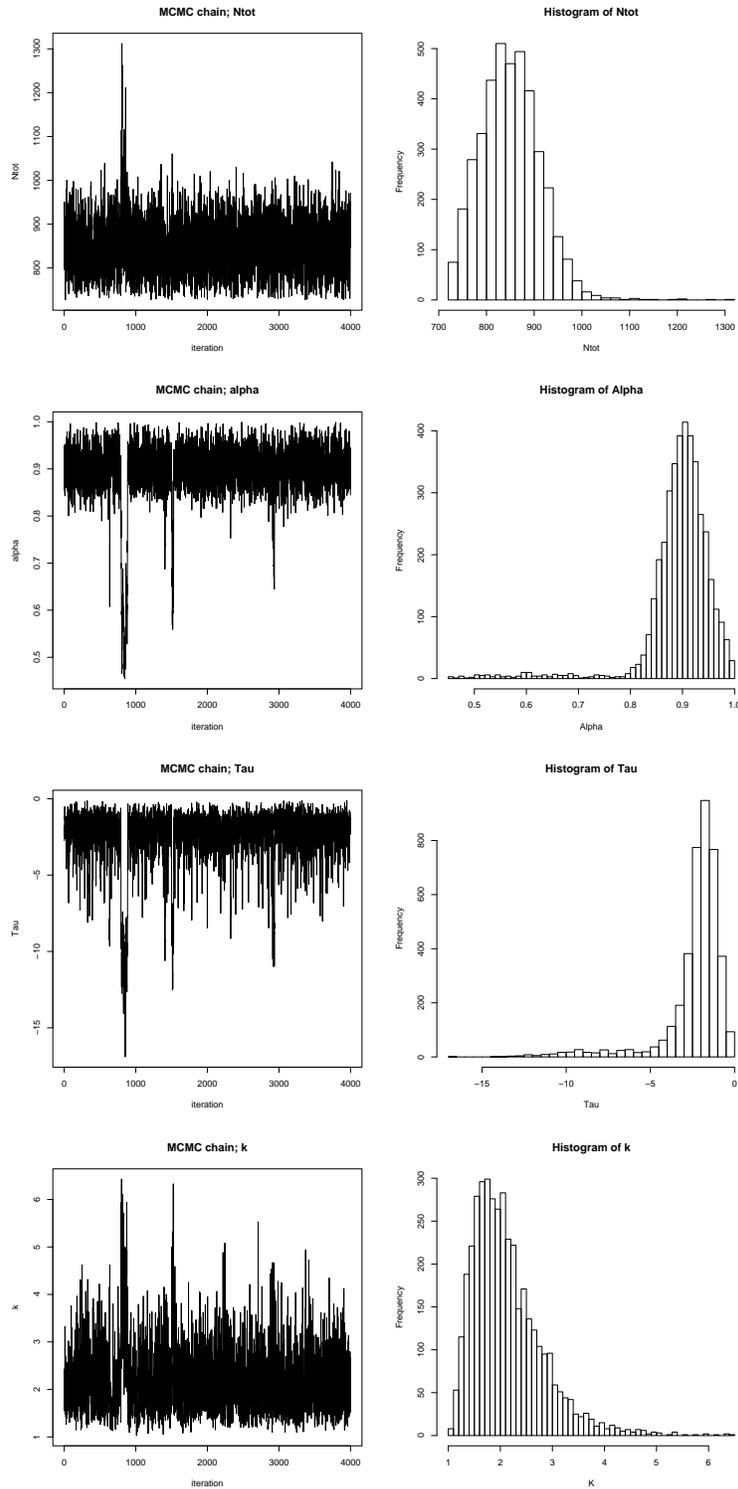


Figure 5.2. The MCMC chains for the outbreak parameters $\{N_{tot}, \alpha, \tau, k\}$, plotted in sequence, from top to bottom, on the left. On the right are the histograms of the samples of the outbreak parameters. These were developed from a time-series 20 days long. The figures for θ were left out for lack of space.

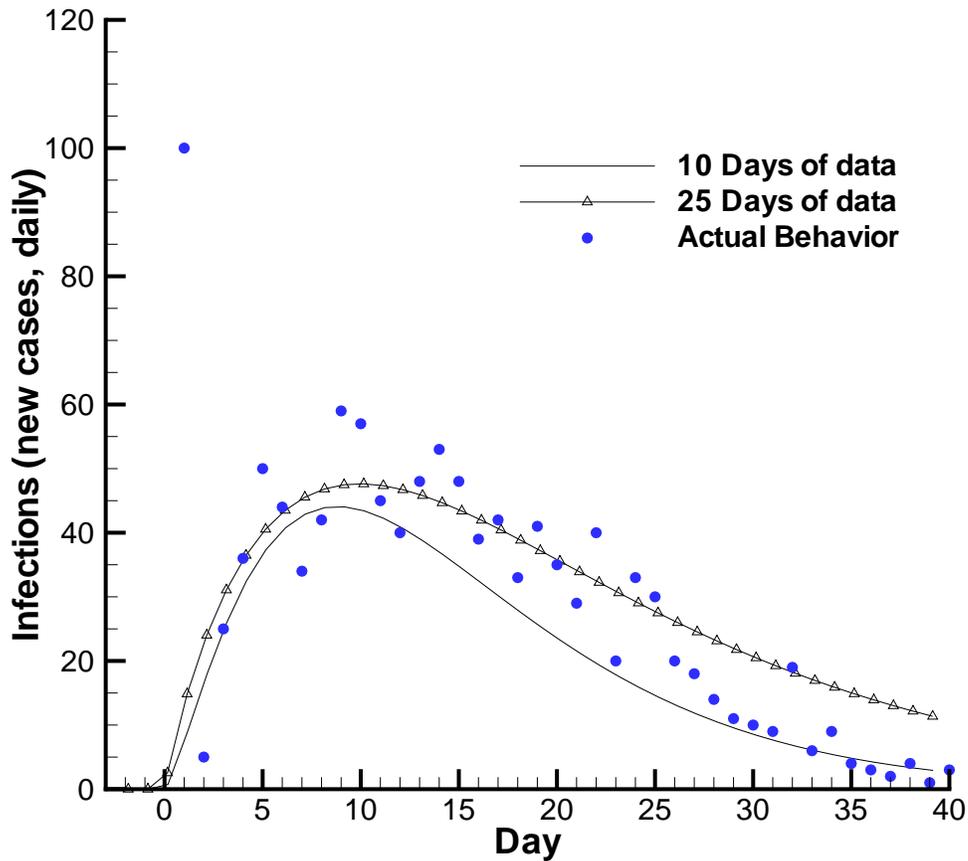


Figure 5.3. The latent infection intensity curve, developed from the maximum likelihood estimates of the outbreak parameters, using 10 and 25 days of data. The true infection intensity is also plotted. The blue dot at the left extreme indicates the index cases; we have not plotted the numerical estimate derived from the time-series.

approximately infer characteristics of an endemic disease (where the infection intensity reaches a constant) by employing a large value for scale parameter θ . However, it is not very useful for a large pandemic that grows in time.

Given the recent interest in swine flu [47, 48], we simulate an influenza epidemic which proves resistant to countermeasures and therefore grows “unbounded” (i.e., before saturation effects become important). This is done with our agent-based simulation capability, with parameters obtained from the 1918 pandemic. These results are then used to infer outbreak parameters.

5.4.1 The outbreak simulation

The disease dynamics for influenza are slightly different from that of plague. While the disease transmission occurs over a social network, the reproductive number does not decay in time (the primary reason why the outbreak can become a pandemic). Transmission is modeled as a Poisson process, but the disease shows 2 extra stages, compared to plague. The different compartments of disease progression are

1. Susceptible (S), people who can be infected.
2. Exposed (E), i.e., people who are incubating the disease, but are not contagious. They move on to the I and A stages described below.
3. Infectious (I), people who are symptomatic and contagious. These people are generally the ones who seek medical help and a fraction of them may be hospitalized and reported.
4. Asymptomatic and contagious (A), this being influenza. In fact most people infected with influenza are in the A category. Such people are not very contagious, but they do exist in large numbers.
5. Recovered (R); people from the I and A stages move into this stage.
6. Dead (D); people from the I stage can move into this stage.

As discussed in [49], only about 36% of the people in the 1918 epidemic in Geneva progressed from the E to the I stage in the (more virulent) “fall” version of the disease; in the “spring” version, only about 10% of the people showed severe symptoms. Further, the asymptomatic patients were far less contagious; the parameter estimates in [49] show that the contagiousness of the asymptomatic were 0.003 and 0.014 times that of the symptomatic patients. In our model, we will assume that the asymptomatic cohort exists, but is not contagious. Mortality rate was 0.7% for the “spring” outbreak and 3.25% for the “fall” outbreak [50]. We also limit our simulation for a short period of time so that births and deaths in the general population do not appreciably affect the progress of the disease.

Each node spends a duration in each of stages E , I and A . These durations are modeled as random variables. Bombardt [7] models the E phase with a lognormal and the I phase with a normal distribution, which we adopt here. Bombardt states the the mean and standard deviation for the E stage are 2 days and 1 day respectively, while those for I stage are 5 days and 1 day. Gani *et al.* [50] find that the mean incubation period (i.e., E stage) is 2 days for pandemic influenza (specifically H5N1 “avian flu”) and 4 days for the symptomatic period. Longini *et al.* [51] and Mills *et al.* [52] cite 1.91 and 4.1 days respectively for the mean values for E and I stages, even though the first publication targets the 1957-1958 and 1968-1969 influenza pandemics while the latter models the 1918 pandemic, with data collected from US cities. The recent swine flu E stage has been estimated to be 1.9 days too [48]. Thus, the characterization of the progress of the disease seems consistent.

Table 5.2. Summary of influenza characterization from a variety of sources. All estimates of duration are in days.

Source	Incubation (E)	Symptomatic (I)	R_0	Remarks
Chowell [49]	2	2	1.49 & 3.75	1918, Geneva
Mills [52]	1.9	4.1	2.0	1918, US cities
Longini [51]	1.9	4.1	1.4	Southeast Asia; “regular” flu
Fraser [48]	1.9	-	1.58	Swine flu; Mexico, 2009
Gani [50]	2.0	4.0	1.39	1957, 1968 pandemics
Bombardt [53]	lognormal; mean = 2; sd = 1	normal; mean = 5; sd = 1		1918 Camps Custer and Valdahon

The basic reproductive number R_0 shows significant variations. Chowell *et al.* [49] report a R_0 of 1.49 for the “spring” outbreak and 3.75 for the “fall” outbreak in Geneva, 1918. US cities, on the other hand, showed an R_0 of 2 in 1918 [52]. Studies for influenza epidemic (mostly H5N1) in Southeast Asia [51, 50] have tended to assume a R_0 of 1.4, which is similar to the R_0 observed in the recent swine-flu epidemic [48] (1.58). We summarize the results reviewed to date in Table 5.2.

For our modeling purposes, we choose a *consensus* set of figures. We will model incubation E as Bombardt does, i.e., as a lognormal distribution with mean 2 days and standard deviation of 1 day. The symptomatic stage I will be modeled using a normal distribution, with a mean of 4 days and standard deviation 1. The asymptomatic stage A will be modeled identical to I , except that the cohort does not suffer deaths. We assume that only 36% of the people coming out of incubation will progress to the I stage, with the rest moving on to the A cohort. Since we aim to capture the 1918 effects, we will assume a mortality rate of 3.25% and a λ that corresponds to a R_0 of 3.75. On recovery, people are assumed to be immune to influenza. These are summarized in Table 5.3 below.

The implementation of the disease model is the same as in Sec. 5.3 and is omitted.

5.4.2 Inference of outbreak parameters

Fig. 5.4 shows the temporal evolution of an outbreak. We start with 100 index cases, who infected 26479 people over the next 42 days. Plotted in red are the new symptomatic cases, collated on a daily basis. Plotted in blue is the latent infection intensity as a function of time. Both the infection intensity and the new symptomatic cases (collated daily) show a monotonic increase.

Table 5.3. Summary of parameters for our influenza model.

Model Parameter	Value
Incubation period (E)	log-normal, mean = 2 days, std. dev. = 1 day
Infectious period (I)	normal, mean = 4 days, std. dev. = 1 day
Asymptomatic period (A)	same as I
Infectious fraction	36 %
Mortality rate	3.25%; applies only to the I cohort
R_0	3.75

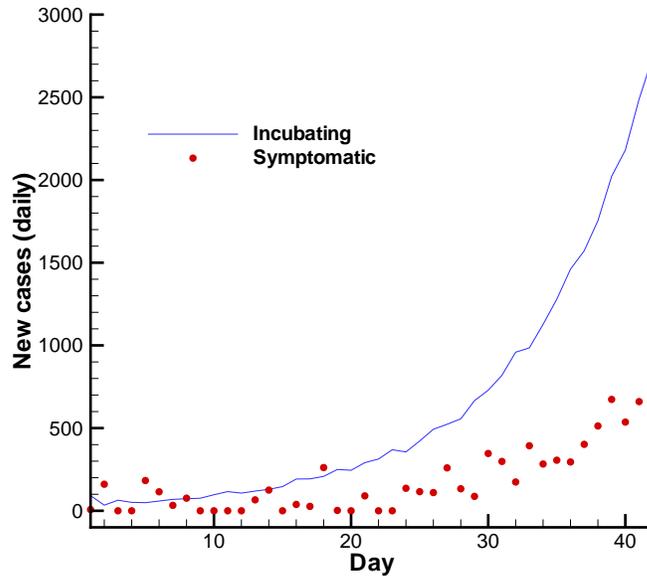


Figure 5.4. Evolution of the influenza epidemic. New cases of symptomatic patients are plotted in red while the latent infection rate (daily infections) are in blue. The simulation was started with 100 index cases distributed in the social network. We see that both the time-series exhibit an upward trend.

Inferring the infection intensity $f(t; \mathbf{p})$ is key to predicting the evolution of the outbreak and consequently the resource requirements. We model the infection intensity as follows

$$f(t; \mathbf{p}) = \frac{a(\exp(t/b) - 1)}{ab(\exp(T/b) - 1) - aT} \quad (5.7)$$

where T is the time duration over which the infection process has occurred. The expression in

Table 5.4. Estimates (medians) of the outbreak parameters developed from time-series of different lengths. The 95% confidence intervals are mentioned in parenthesis and the true values are in square brackets.

Variable	m = 15	m = 25
N_{tot}	1671 (795, 13873) [1474]	5220 (2994, 29691) [4614]
α	0.54 (0., 0.92) [0.93]	0.78 (0.04, 0.98) [0.98]
τ	-4 (-19., -0.1) [-1]	-6.8 (-28, -0.27) [-1]
a	2.73 (1.06, 15.8)	2.8 (1.2, 14.8)
b	104.9 (26, 1.2×10^5)	192 (11.3, 2.1×10^5)

Eq. 5.7 obeys the normalization Eq. 5.1. a and b are the governing parameters of the infection intensity and form the objects of inference from data. Following the description in Sec. 5.3, we perform an inference of the outbreak parameters using a time-series of different lengths. The problem was reparameterized in terms of the logarithms of the quantities being inferred, i.e., $\{N_{tot}, \alpha, \tau, a, b\}$. The priors used are

$$\begin{aligned} \log(N_{tot}) &\sim N(\log(10^3), 10), \\ \alpha &\sim \mathcal{B}(1.25, 1.25), \\ \log(-\tau) &\sim N(0, 1), \\ \log(a) &\sim N(0, 5^2), \\ \log(b) &\sim N(0, 5^2) \end{aligned}$$

For the purposes of this study, the standard deviation of the measurement error σ is set to 50.

In Table 5.4, we summarize the outbreak parameter estimates drawn from time-series of different lengths. We see that the size of the epidemic N_{tot} is estimated fairly accurately, but there are significant errors in estimates of the rest of the parameters. This is conjectured to be due to the exponential nature of the infection intensity, which makes the predictions n_i^{model} very sensitive to the model parameters. We are currently investigating how reparameterizing may ameliorate the sensitivity and allow more robust inferences.

This page intentionally left blank.

Chapter 6

Conclusions

Our study of resource allocation techniques in reload scenarios was driven by the realization that mounting a quick and efficient response to a bioattack holds the largest potential to reduce casualties and minimize impact on the affected population. To date, early warning via detection of aerosolized pathogens by environmental sensors has been viewed as the optimal way of determining when and how to mount a medical response. We consider the case when an attack may be not be detected by such sensors either because the site was not instrumented or if the pathogen was introduced via a vector (which, for communicable diseases, could be humans). In such a case, the estimation of resource allocation have to be performed using the time-series of morbidity (diagnosed cases etc) that would result from the ensuing outbreak. In our study, we have restricted ourselves to resources like medical equipment and personnel which are difficult to gather and transport.

We have developed an approach that allows the estimation of attack/outbreak parameters from short time-series of morbidity data. The attack parameters, viz., the number of infected people, the time of attack and the dose, can be related to the time-dependent demand for medical resources directly using existing epidemic and resource-use models. The attack parameters are estimated probabilistically, which reflects the uncertainty due to lack/quality of the data. The attack parameters are used to bound the possible realizations of the resource demand; thereafter, a stochastic optimization algorithm develops a resource allocation profile (in time). Resources are assumed to be insufficient (hence an efficient allocation is paramount). This can be used to dispatch resources in the short term and plan for transportation needs in the long term. We find that our resource allocations amount to hedging – they render the probability density function of casualties resulting from a resource allocation profile compactly supported, indicating a significant reduction of the probability of an exceedingly bad outcome. The price for reducing this probability is the increased probability (almost a certainty) of a smaller, perhaps acceptable, level of casualties. Our test cases have involved single-site attacks as well as staggered attacks on multiple sites, where the resource demands at different sites have different levels of uncertainty. The observations above hold true in both cases; further, for multi-site attacks, our allocation technique ensures “fairness”, i.e., there are no hot-spots of risk which could endanger the entire system via cascading failures. We observe, empirically, that resource allocation profile does not vary significantly (by more than 25%) after about 6 days of data. This lack of volatility in the allocation profile is helpful since it assists in planning for the mobilization of infrastructural capabilities (and to some extent, resource reserves).

Our algorithm involves a “free” parameter. An optimal value of this “free” parameter cannot

be gauged from the time-series of morbidity and thus is exogenous to the problem at hand. It represents a “risk appetite” and takes the form of a constraint on the daily allocation of resources. This constraint can be used to prevent an over-allocation of resources, a recourse that may prove important in the aftermath of an attack, if further attacks are expected and resources have to be husbanded. More practically, this constraint may be used to enforce resource allocation which conform to the dictates of the transportation infrastructure. In the immediate aftermath of an attack, the available rolling/transportation stock may simply be insufficient. The constraint can be made time-dependent, to reflect the mobilization of infrastructural capabilities to meet an emergency.

The technique works well for outbreaks/attacks carried out with pathogens causing non-communicable diseases. The problem of epidemics caused by communicable diseases is more difficult, primarily as it poses a harder (non-stationary) estimation problem. We have outlined an estimation approach here, but it is limited by the simple epidemic model employed in the estimation algorithm. This is an area that should be investigated in more detail; our resource allocation techniques can find use in responding to endemic diseases (and their possible pandemic variants). Our example in Chp. 5 used influenza modeled on the fall variant of the 1918 pandemic.

In our study, we have not attempted to model the imperfections of transportation infrastructure beyond assuming that there is an initial mobilization delay. However a 25% volatility in resource allocation is not trivial, and advanced routing algorithms (for rolling stock) may be required to accommodate it. Routing on networks is an enduring problem in operations research, and our technique for estimating uncertain demands/throughputs at certain points in the network may allow the application of network routing algorithms to a new field of national security/interest.

References

- [1] J-E. C. Holty, D. M. Bravata, H. Liu, R. A. Olshen, K. M. McDonald, and D. K. Owens. Systematic review: A century of inhalational anthrax cases from 1900 to 2005. *Annals of Internal Medicine*, 144:270–280, 2006.
- [2] R. Danzig. *Catastrophic bioterrorism: What is to be done?* Bernan Press, PA, USA, 2003.
- [3] D. M. Edwards, T. H. West, S. Gordon, I. Chumfong, and N. J. Gleason. Biological reference scenarios. SAND Report SAND2006-3837, Sandia National Laboratories, Livermore, CA 94551-0969, July 2006. Official Use Only.
- [4] J. Walden and E. H. Kaplan. Estimating time and size of bioterror attack. *Emerging Infectious Diseases*, 10(7):1202–1205, 2004.
- [5] Matthew Meselson, Jeanne Guillemin, Martin Hugh-Jones, Alexander Langmuir, Ilone Popova, Alexis Shelokov, and Olga Yampolskaya. The Sverdlovsk anthrax outbreak of 1979. *Science*, 266:1202–1208, 1994.
- [6] R. Brookmeyer, N. Blades, M. Hugh-Jones, and D. A. Henderson. The statistical analysis of truncated data: application to the Sverdlovsk anthrax outbreak. *Biostatistics*, 2:233–247, 2001.
- [7] Ron Brookmeyer and Natalie Blades. Statistical models and bioterrorism: Application to the U.S. anthrax attacks. *Journal of the American Statistical Association*, 98(464):781–788, 2003.
- [8] John A. Jernigan et al. Bioterrorism-related inhalational anthrax: The first 10 cases reported in the United States. *Emerging Infectious Diseases*, 7(6):933–944, 2001.
- [9] William R. Hogan, Gregory F. Cooper, Garrick L. Wallstrom, Michael M. Wagner, and Jean-Marc Depinay. The Bayesian aerosol release detector: An algorithm for detecting and characterizing outbreaks caused by an atmospheric release of *Bacillus anthracis*. *Statistics in Medicine*, 26:5225–5252, 2007.
- [10] F.-C. Tsui, J. U. Espino, V. M. Dato, P. H. Gesteland, J. Hutman, and M. M. Wagner. Technical Description of RODS: A Real-time Public Health Surveillance System. *Journal of the American Medical Informatics Association*, 10(5):399–408, 2003.
- [11] D. B. Turner. *Workbook of Atmospheric Dispersion Estimates: An Introduction to Dispersion Modeling*. Lewis Publishers, CRC Press LLC, 2000 N.W. Corporate Blvd. Boca Raton, FL 33431, 1994.
- [12] H. N. Glassman. Discussion on industrial inhalational anthrax. *Bacteriological Review*, 30:657–659, 1966.

- [13] Judith Legrand, Joseph R. Egan, Ian M. Hall, Simon Cauchemez, Steve Leach, and Neil M. Ferguson. Estimating the location and timing of a covert anthrax release. *Public Library of Science - Computational Biology*, 5(1), 2009.
- [14] J. Ray, Y. M. Marzouk, M. Kraus, and P. Fast. A Bayesian method for characterizing distributed micro-releases: II. Inference under model uncertainty with short time-series data. SAND Report SAND2006-7568, Sandia National Laboratories, Livermore, CA 94551-0969, December 2006. Unclassified unlimited release; available at <http://www.caip.rutgers.edu/~jaray/Inverse.html>.
- [15] W. R. Hogan and G. L. Wallstrom. Approximating the sum of lognormal distributions to enhance models of inhalational anthrax. In *Quantitative Methods in Defense and National Security*, Fairfax, VA, 2007. Available at: <http://www.galaxy.gmu.edu/QMDNS2007/QMDNS2007-booklet.pdf> (Accessed on 30 March 2007).
- [16] P. D. O’Neill and G. O. Roberts. Bayesian inference of partially observed stochastic epidemics. *Journal of the Royal Statistical Society, Series A*, 162:121–129, 1999.
- [17] Martin Eichner and Klaus Dietz. Transmission potential of smallpox: Estimates based on detailed data from an outbreak. *American Journal of Epidemiology*, 158(2):110–117, 2003.
- [18] Yang Yang, M. Elizabeth Halloran, Johathan D. Sugimota, and Ira M. Longini. Detecting human-to-human transmission of avian influenza A (H5N1). *Emerging Infectious Diseases*, 13(9):1348–1353, 2007.
- [19] Luis M. A. Bettencourt and Ruy M. Ribeiro. Real-time Bayesian estimation of the epidemic potential of emerging infectious diseases. *Public Library of Science – One*, 3(5), 2008.
- [20] Gerardo Chowell, Hiroshi Nishiura, and Luis M. A. Bettencourt. Comparative estimation of the reproduction number for pandemic influenza from daily case notification data. *Journal of the Royal Society – Interface*, 4, 2007.
- [21] S. Cauchemez and F. Carrat, C. Viboud, A. J. Valleron, and P. Y. Boelle. A Bayesian MCMC approach to study transmission of influenza: application to household longitudinal data. *Statistics in Medicine*, 23:3469–3487, 2004.
- [22] Tom Britton and P. O’Neill. Bayesian inference for stochastic epidemics in populations with random social structure. *Scandinavian Journal of Statistics*, 29:375–390, 2002.
- [23] J. Ray, B. M. Adams, K. D. Devine, Y. M. Marzouk, M. M. Wolf, and H. N. Najm. Distributed micro-releases of bioterror pathogens - Threat characterizations and epidemiology from uncertain patient observables. Technical Report SAND2008-6044, Sandia National Laboratories, Livermore, CA 94551-0969, October 2008. Unclassified, unlimited release.
- [24] R. Brookmeyer and M. H. Gail. A method for obtaining short-term projections and lower bounds on the size of the AIDS epidemic. *Journal of the American Statistical Association*, 83(402):301–308, 1988.

- [25] Paul T. Boggs, David M. Gay, and Jaideep Ray. Probabilistic attack reconstruction and resource estimation in “reload” scenarios. In *Conference Publication for 2009 IEEE International Conference on Technologies for Homeland Security*, 2009. <http://ieeehomelandsecurityconference.org/>.
- [26] D. W. Pierce. *The MIT Dictionary of Modern Economics*. The MacMillan Press, 1986.
- [27] R. G. Palmer, W. B. Arthur, J. H. Holland, B. LeBaron, and P. Tayler. Artificial economic life: A simple model of a stockmarket. *Physica D*, 75:264–274, 1994.
- [28] R. Axelrod. *The complexity of cooperation : Agent-based models of competition and collaboration*. Princeton University Press, Princeton, NJ, 1997.
- [29] W. B. Arthur, J. H. Holland, B. Lebaron, R. G. palmer, and P. Tayler. The economy as a complex evolving system II. In w. B. Arthur, S. Durlauf, and D. Lane, editors, *Santa Fe Institute Studies in the Sciences of Complexity*, pages 15–42. Addison–Wesley, Reading, MA, 2007.
- [30] J. Epstein and R. Axtell. *Growing artifical societies: Social science from the bottom up*. MIT Press, Cambridge, MA, 1996.
- [31] A Weidlich and D. Veit. A critical survey of agent-based wholesale electricity market models. *Energy Economics*, 30(4):1728–1759, 2008.
- [32] E. Bonebeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(3):7280–7287, 2002.
- [33] M. A. Ehlen, A. J. Scholand, and K. L. Stamber. The effects if residential real-time pricing contracts on transco loads, pricing and profitability: Simulations using the N-ABLE agent-based model. *Energy Economics*, 29(2):211–227, 2007.
- [34] J. A. Sprigg and M. A. Ehlen. Comparative dynamics in an overlapping-generations model: Tthe effects of quasi-rational discrete choice on finding and maintaining nash equilibrium. *Computational Economics*, 29(1):69–96, 2007.
- [35] N. Basu, R. Pryor, and T. Quint. ASPEN: A microsimulation model of the economy. *Computational Economics*, 12(3):223–241, 1998.
- [36] E. D. Eidson and M. A. Ehlen. Nisac agent-based laboratory for economics (n-able): Overview of agent and simulation architectures. SAND Report SAND2005-0263, Sandia National Laboratories, Albuquerque, NM, February 2005. Unclassified unlimited release.
- [37] J. Ray, P. T. Boggs, D. M. Gay, Y. M. Marzouk, and H. N. Najm. A Bayesian approach for estimating bioterror attacks from patient data. *Statistics in Medicine*, 2009. Under review.
- [38] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Chapman and Hall, 1996.
- [39] CRAN site for the mcgibbist package. <http://cran.r-project.org/web/packages/mcgibbsit/index.html>.

- [40] R. Fourer, D. M. Gay, and B. W. Kernighan. A modeling language for mathematical programming. *Management Science*, 36(5):519–554, 1990.
- [41] Robert Fourer, David M. Gay, and Brian W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. Duxbury Press/Brooks/Cole Publishing Co., second edition, 2003.
- [42] <http://www.ampl.com/>.
- [43] <http://www.ilog.com/products/cplex/>.
- [44] Stephen Eubank, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltan Toroczka, and Nan Wang. Modelling disease outbreaks in realistic urban social networks. *Nature*, 429:180–184, 2004.
- [45] R. Gani and S. Leach. Epidemiological determinants for modeling pneumonic plague outbreaks. *Emerging Infectious Diseases*, 10(4):608–614, 2004.
- [46] H. Nishiura, M. Schwehm, M. Kakehashi, and M. Eichner. Transmission potential of primary plague: time inhomogeneous evaluation based on historical documents of the transmission network. *Journal of Epidemiology and Community Health*, 60:640–645, 2006.
- [47] CDC 2009 H1N1 Swine Flu. http://www.cdc.gov/h1n1flu/general_info.htm.
- [48] Christophe Fraser et al. Pandemic potential of a strain of influenza A (H1N1): Early findings. *Science*, 2009. To appear; download from <http://www.sciencexpress.org>; DOI:10.1126/science.1176062.
- [49] G. Chowell, C. E. Ammon, N. W. Hengartner, and J. M. Hyman. Transmission dynamics of the reat influenza pandemic of 1918 in geneva, switzerland: Assessing the effects of hypothetical interventions. *Journal of Theoretical Biology*, 241(2):193–204, 2006.
- [50] R. Gani, H. Hughes, D. Fleming, T. Griffin, J. Medlock, and S. Leach. Potential impact of antiviral drug use during influenza pandemic. *Emerging Infectious Diseases*, 11(9):1355–1362, 2005.
- [51] I. M. Longini, A. Nizam, S. Xu, K. Ungchusak, W. Hanshaoworakul, D. .A. T. Cummings, and M. E. Halloran. Containing pandemic influenza at source. *Science*, 309:1083–1087, 2005.
- [52] C. E. Mills, J. M. Robbins, and M. Lipsitch. Transmissibility of 1918 pandemic influenza. *Nature*, 432:904–906, 2004.
- [53] J. N. Bombardt Jr and H. E. Brown. Potential influenza effects on military populations. Technical Report P-3786, Institute for Defense Analyses, December 2003.

DISTRIBUTION:

2	MS 9159	Jaideep Ray, 08964
1	MS 9159	Paul T. Boggs, 08961
1	MS 1318	David M. Gay, 01411
1	MS 9158	Michelle N. Lemaster, 08961
1	MS 1138	Mark A. Ehlen, 06321
1	MS 0899	Technical Library, 8944 (electronic)
1	MS 0123	D. Chavez, LDRD Office, 1011

This page intentionally left blank.

